


IBM Flex System V7000 Storage Node

Introduction and Implementation Guide

Introduction to IBM Flex System family,
features, and functions

IBM Flex System V7000 Storage Node
hardware overview

Host configuration guide



Sangam Racherla
Eva Ho
Carsten Larsen
Kim Serup
John Sexton
Mansoor Syed
Alexander Watson

Redbooks



International Technical Support Organization

**IBM Flex System V7000 Storage Node
Introduction and Implementation Guide**

March 2013

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

First Edition (March 2013)

This edition applies to IBM Flex System V7000 Storage Node Version 6.4.

Note: This book is based on a pre-GA version of a product and some of the features described in this book might change when the product becomes generally available. We suggest that you consult the product documentation or follow-on versions of this publication for more current information.

© Copyright International Business Machines Corporation 2013. 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	xi
Trademarks	xii
Preface	xiii
The team who wrote this book	xiii
Now you can become a published author, too!	xv
Comments welcome	xv
Stay connected to IBM Redbooks	xvi
Chapter 1. Introduction to IBM Flex Systems and IBM PureSystems offerings	1
1.1 IBM PureSystems overview	2
1.1.1 Product names	3
1.1.2 IBM PureFlex System	4
1.1.3 IBM PureApplication System	6
1.2 IBM PureFlex System building blocks	8
1.2.1 Highlights	8
1.2.2 Components	9
1.3 IBM Flex System Enterprise Chassis	10
1.3.1 Chassis power supplies	13
1.3.2 Fan modules and cooling	14
1.4 Compute nodes	15
1.4.1 IBM Flex System x440 Compute Node	15
1.4.2 IBM Flex System x240 Compute Node	16
1.4.3 IBM Flex System x220 Compute Node	19
1.4.4 IBM Flex System p260 and p24L Compute Nodes	20
1.4.5 IBM Flex System p460 Compute Node	22
1.5 I/O modules	24
1.5.1 IBM Flex System Fabric CN4093 10 Gb Converged Scalable Switch	25
1.5.2 IBM Flex System Fabric EN4093 and EN4093R 10 Gb Scalable Switch	26
1.5.3 IBM Flex System EN4091 10 Gb Ethernet Pass-thru	27
1.5.4 IBM Flex System EN2092 1 Gb Ethernet Scalable Switch	27
1.5.5 IBM Flex System FC5022 16 Gb SAN Scalable Switch	28
1.5.6 IBM Flex System FC3171 8 Gb SAN Switch	28
1.5.7 IBM Flex System FC3171 8 Gb SAN Pass-thru	29
1.5.8 IBM Flex System IB6131 InfiniBand Switch	29
1.6 Introduction to IBM Flex System storage	30
1.6.1 IBM Storwize V7000 and IBM Flex System V7000 Storage Node	30
1.6.2 Benefits and value proposition	34
1.6.3 Data Protection features	34
1.7 External storage	35
1.7.1 Storage products	35
1.7.2 IBM Storwize V7000	36
1.8 Useful websites	36
Chapter 2. Introduction to IBM Flex System V7000 Storage Node	39
2.1 IBM Flex System V7000 Storage Node overview	40
2.2 IBM Flex System V7000 Storage Node terminology	41
2.3 IBM Flex System V7000 Storage Node	43
2.3.1 IBM Flex System V7000 Storage Node capabilities	44

2.3.2	IBM Flex System V7000 Storage Node functions	44
2.3.3	IBM Flex System V7000 Storage Node licensing	46
2.4	IBM Flex System V7000 Storage Node hardware	47
2.4.1	Control canister	48
2.4.2	Expansion canister	52
2.4.3	Supported disk drives	53
2.4.4	IBM Storwize V7000 expansion enclosure	54
2.4.5	SAS cabling requirements	55
2.5	IBM Flex System V7000 Storage Node components	56
2.5.1	Hosts	57
2.5.2	Control canisters	57
2.5.3	I/O groups	57
2.5.4	Clustered system	58
2.5.5	RAID	59
2.5.6	Managed disks	60
2.5.7	Quorum disks	60
2.5.8	Storage pools	61
2.5.9	Volumes	64
2.5.10	Thin-provisioned volumes	66
2.5.11	Mirrored volumes	67
2.5.12	Easy Tier	68
2.5.13	Real-time Compression (RTC)	69
2.5.14	iSCSI	69
2.5.15	Fibre Channel over Ethernet (FCoE)	71
2.6	Advanced copy services	72
2.6.1	FlashCopy	72
2.6.2	IBM Flex System V7000 Remote Mirroring software	73
2.6.3	Remote Copy licensing	74
2.6.4	Synchronous / Asynchronous Remote Copy	74
2.6.5	Copy Services configuration limits	75
2.7	Management and support tools	76
2.7.1	IBM Assist On-site and remote service	76
2.7.2	Event notifications	77
2.7.3	SNMP traps	77
2.7.4	Syslog messages	77
2.7.5	Call Home email	78
2.8	Useful references from Storwize V7000 websites	78
2.9	IBM virtual storage learning videos on YouTube	80
	Chapter 3. Systems management	81
3.1	System Management overview	82
3.1.1	Integrated platform management tools	82
3.1.2	IBM Flex System storage management	84
3.1.3	Storage management interfaces	85
3.2	IBM Flex System Chassis Management Module (CMM)	88
3.2.1	Overview of IBM Flex System Chassis Management Module	88
3.2.2	Accessing the CMM	90
3.2.3	Viewing and configuring IP addresses of chassis components	92
3.2.4	Accessing I/O modules using CMM	94
3.2.5	Managing storage using IBM Flex System Chassis Management Module	96
3.2.6	Data collection using CMM	106
3.3	Flex System Manager (FSM)	110
3.3.1	Overview of IBM Flex System Manager (FSM)	110

3.3.2	IBM Flex System Manager storage management features	112
3.3.3	Logging in to the IBM Flex System Manager Node	116
3.3.4	Overview of IBM Flex System Manager and IBM FSM Explorer	117
3.3.5	Accessing I/O modules using FSM	130
3.3.6	Data collection using FSM	135
3.3.7	Managing storage using IBM Flex System Manager	139
Chapter 4.	IBM Flex System V7000 Storage Node initial configuration	149
4.1	Planning overview	150
4.1.1	Hardware planning	150
4.1.2	SAN configuration planning	151
4.1.3	LAN configuration planning	151
4.1.4	Management IP address considerations	152
4.1.5	Service IP address considerations	152
4.1.6	Management interface planning	153
4.2	Initial setup for IBM Flex System V7000 Storage Node	153
4.2.1	Using FSM for initial setup	155
4.2.2	Using CMM for initial setup	159
4.3	IBM Flex System V7000 Storage Node Setup Wizard	162
4.4	System management	167
4.4.1	Graphical User Interface (GUI)	168
4.4.2	Launching IBM Flex System V7000 Storage Node GUI from CMM	170
4.5	Service Assistant	171
4.5.1	Changing the Service IP address	171
4.6	Command-Line interface (CLI)	173
4.7	Recording system access information	174
Chapter 5.	IBM Flex System V7000 Storage Node GUI interface	175
5.1	Overview of IBM Flex System V7000 Storage Node management software	176
5.1.1	Access to the Graphical User Interface	176
5.1.2	Graphical User Interface layout	177
5.1.3	Navigation	178
5.1.4	Multiple selections	181
5.1.5	Status Indicators menus	182
5.2	Home menu	184
5.3	Monitoring menu	185
5.3.1	Monitoring System Details menu	185
5.3.2	Monitoring Events menu	190
5.3.3	Monitoring Performance menu	198
5.4	Pools menu	200
5.4.1	Volumes by Pool menu	201
5.4.2	Internal Storage menu	202
5.4.3	External Storage menu	203
5.4.4	System Migration tools	203
5.4.5	MDisks by Pools menu	204
5.5	Volumes menu	207
5.5.1	The Volumes menu	208
5.5.2	Volumes by Pool menu	210
5.5.3	Volumes by Host menu	212
5.6	Hosts menu	214
5.6.1	Hosts menu	215
5.6.2	Ports by Host menu	217
5.6.3	Host Mappings menu	219

5.6.4	Volumes by Host menu	221
5.7	Copy Services menu	221
5.7.1	FlashCopy menu	222
5.7.2	FlashCopy Consistency Group menu	224
5.7.3	FlashCopy Mapping menu	225
5.7.4	Remote Copy and the Partnerships menu	225
5.8	Access menu	226
5.8.1	Users menu	227
5.8.2	Audit Log menu	229
5.9	Settings menu	230
5.9.1	Event Notification menu	231
5.9.2	Directory Services	232
5.9.3	Network menu	233
5.9.4	Support menu	235
5.9.5	General menu	240
Chapter 6.	Basic volume and host configuration.	241
6.1	Storage provisioning from the IBM Flex System V7000 Storage Node.	242
6.1.1	Creating a generic volume	244
6.1.2	Creating a thin-provisioned volume.	245
6.1.3	Creating a mirrored volume.	248
6.1.4	Creating a thin-mirror volume	251
6.1.5	IBM Real-time Compression	254
6.2	Creating a new host	256
6.2.1	Creating a Fibre Channel attached host	257
6.2.2	Creating an iSCSI attached host.	259
6.3	Mapping a volume to the host.	262
6.3.1	Mapping newly created volumes to the host using the wizard	262
6.3.2	Additional features	264
Chapter 7.	Storage Migration Wizard	265
7.1	Preparing for data migration	266
7.2	Migrating the data using the Storage Migration Wizard	266
7.2.1	Checking the Windows 2008 host before upgrade	267
7.2.2	Remapping the disk to the IBM Flex System V7000 Storage Node	267
7.2.3	Storage Migration Wizard on the IBM Flex System V7000 Storage Node	271
7.2.4	Checking that the Windows server can see its disks	284
7.2.5	Finalizing the migration	287
7.2.6	Mapping to host performed after the migration has begun	289
7.2.7	Renaming the volume	292
Chapter 8.	Storage pools	295
8.1	Working with internal drives	296
8.1.1	Actions on internal drives	300
8.1.2	Configuring internal storage	303
8.2	Working with MDisks	313
8.2.1	Adding MDisks to storage pools	316
8.2.2	Importing MDisks	320
8.2.3	RAID action for MDisks.	329
8.2.4	Selecting the tier for MDisks	331
8.2.5	Additional actions on MDisks	333
8.2.6	Properties for Mdisks	335
8.3	Working with storage pools	337

Chapter 9. IBM Flex System V7000 Storage Node Copy Services	341
9.1 Services provided	342
9.2 FlashCopy	342
9.2.1 Business requirements for FlashCopy	342
9.2.2 FlashCopy functional overview	343
9.2.3 Planning for FlashCopy	351
9.2.4 Managing FlashCopy using the GUI	353
9.2.5 Managing FlashCopy mapping	359
9.3 Remote Copy	380
9.3.1 Remote Copy concepts	380
9.3.2 Global Mirror with Change Volumes	387
9.3.3 Remote Copy planning	392
9.4 Troubleshooting Remote Copy	396
9.4.1 1920 error	397
9.4.2 1720 error	399
9.5 Managing Remote Copy using the GUI	399
9.5.1 Managing cluster partnerships	399
9.5.2 Deleting a partnership	405
9.5.3 Managing a Remote Copy consistency group	415
Chapter 10. Volume mirroring and migration	423
10.1 Volume mirroring and migration options	424
10.2 Tunable timeout	424
10.3 Usage of mirroring for migration	424
10.4 Managing Volume Mirror and migration with the GUI	425
Chapter 11. SAN connections and configuration	431
11.1 Storage Area Network overview	432
11.2 Connection to chassis I/O modules	433
11.2.1 I/O module configurations	433
11.2.2 I/O module connection summary	434
11.3 iSCSI connectivity	435
11.3.1 Session establishment and management	436
11.3.2 iSCSI initiators	436
11.3.3 iSCSI multisection configuration and support	438
11.3.4 iSCSI multipath connectivity	439
11.3.5 Configuring multiple iSCSI host links	440
11.4 FCoE connectivity	442
11.4.1 FCoE protocol stack	443
11.4.2 Converged Network Adapters	444
11.4.3 FCoE port types	445
11.4.4 Configuring CN4093 for FCoE connectivity	447
11.5 Fibre Channel connectivity	455
11.5.1 The concept of layers	456
11.5.2 Fibre Channel topologies	458
11.5.3 FC addressing and port types	459
11.5.4 Zoning a compute node for storage allocation	460
11.5.5 Multipathing	468
11.5.6 FC Switch Transparent Mode and N_Port ID Virtualization (NPIV)	469
11.6 Storage Area Network summary	471
Chapter 12. Host configuration	473
12.1 Host configuration overview	474
12.2 Discovering volumes from the host and multipath settings	475

12.3	Windows host attachment	476
12.3.1	Windows 2008 R2 Fibre Channel volume attachment	476
12.3.2	Windows 2008 R2 iSCSI volume attachment	480
12.3.3	Windows 2008 R2 FCoE volume attachment	493
12.4	VMware ESX host attachment	498
12.4.1	VMware ESX Fibre Channel attachment	498
12.4.2	VMware ESX iSCSI attachment	504
12.5	AIX host attachment	512
12.5.1	Configuring the AIX compute node for FC connectivity	512
12.5.2	Operating system versions and maintenance levels	513
12.5.3	Checking connectivity to the IBM Flex System V7000 Storage Node	514
12.5.4	Installing the 2145 host attachment support package	515
12.5.5	Subsystem Device Driver Path Control Module	516
12.6	Linux host attachment	520
12.6.1	Linux Fibre Channel attachment	520
12.6.2	Linux distributions	520
Chapter 13.	Maintenance and troubleshooting	521
13.1	Reliability, availability, and serviceability (RAS)	522
13.2	Hardware and LED descriptions	522
13.2.1	Understanding the system state using the control enclosure LEDs	522
13.2.2	Understanding the system state using the expansion enclosure LEDs	528
13.2.3	Power-on self-test	531
13.2.4	Powering on using LED indicators	531
13.3	Monitoring system status and health	532
13.3.1	Using FSM for status and health	532
13.3.2	System Status and Health tasks	534
13.4	Managing storage nodes	536
13.4.1	Using FSM Chassis Manager page to manage storage node	536
13.4.2	Using FSM Storage Management to manage storage node	537
13.5	Configuration backup and restore process	538
13.6	Software upgrade	539
13.6.1	Choosing an upgrade method	539
13.6.2	Upgrading the system software	539
13.7	Drive firmware upgrade	545
13.8	Troubleshooting	546
13.8.1	Using the CMM for troubleshooting tasks	546
13.8.2	Using the management GUI for troubleshooting tasks	549
13.8.3	Removing and replacing parts for troubleshooting and resolving problems	550
13.8.4	Event reporting	551
13.8.5	Viewing the event log	551
13.8.6	Error event IDs and error codes	552
13.9	Audit log navigation	555
13.10	Support data collection	555
13.10.1	Collecting System Management Server service data using the CMM	557
13.10.2	Collecting support files using FSM	558
13.11	Using event notifications	559
13.12	Configuring Call Home	560
13.12.1	Configuring Call Home if FSM is not included	560
13.12.2	Configuring Call Home if FSM is included	562
13.13	IBM Flex System V7000 Storage Node power on and off	567
13.13.1	Powering on the system	567
13.13.2	Powering off the system	568

13.13.3 Shutting down using IBM Flex System V7000 command-line interface	568
13.13.4 Powering off a node using the CMM.	569
13.13.5 Powering on a node using the CMM.	569
Appendix A. CLI setup and configuration.	571
Command-line interface	572
Basic setup	572
Example commands	581
Related publications	587
IBM Redbooks publications	587
Other publications	587
Online resources	588
Help from IBM	589

Notices

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

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

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

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

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

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

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

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

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


COPYRIGHT LICENSE:

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

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.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:

Active Cloud Engine™	IBM Flex System™	Redbooks (logo)  ®
Active Memory™	IBM Flex System Manager™	ServerProven®
AIX®	IBM®	Storwize®
BladeCenter®	Power Systems™	System i®
DS4000®	POWER7®	System p®
DS6000™	PowerVM®	System Storage DS®
DS8000®	POWER®	System Storage®
Easy Tier®	PureApplication™	System x®
Electronic Service Agent™	PureFlex™	Tivoli®
EnergyScale™	Real-time Compression™	X-Architecture®
FlashCopy®	Redbooks®	XIV®
HACMP™	Redpapers™	xSeries®

The following terms are trademarks of other companies:

Intel Xeon, Intel, Pentium, 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.

Linux is a trademark of Linus Torvalds 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.

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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

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

Preface

IBM® Flex System™ products are ideally suited for data center environments that require flexible, cost-effective, secure, and energy-efficient hardware. The IBM Flex System V7000 Storage Node is the latest addition to the IBM Flex Systems product family and is a modular storage system designed to fit into the IBM Flex System Enterprise chassis.

When purchased in the IBM PureFlex™ configurations, the IBM Flex System V7000 Storage Node is configured from the factory into the hardware solution purchased. If, however, the configuration wanted is not offered in the predefined offerings, then a “Build to Order” configuration is designed to meet your needs.

The IBM Flex System V7000 Storage Node includes the capability to virtualize its own internal storage in the same manner as the IBM Storwize® V7000 does. It is designed to be a scalable internal storage system to support the compute nodes of the IBM Flex System environment.

This IBM® Redbooks® publication introduces the features and functions of the IBM Flex System V7000 Storage Node through several examples. This book is aimed at pre-sales and post-sales technical support and marketing personnel and storage administrators. It can help you understand the architecture of the IBM Flex System V7000 Storage Node, how to implement it, and how to take advantage of the industry leading functions and features.

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, Raleigh Center.

Sangam Racherla is an IT Specialist and Project Leader working at the International Technical Support Organization (ITSO), San Jose Center. He holds a degree in Electronics and Communication Engineering and has eleven years of experience in the IT field. Sangam has been with the ITSO for the past nine years and has extensive experience installing and supporting the ITSO lab equipment for various IBM Redbooks publications projects. His areas of expertise include Microsoft Windows, Linux, IBM AIX®, IBM System x®, and IBM System p® servers, and various SAN and storage products.

Eva Ho is a worldwide Product Engineer support for IBM Flex System V7000 with the IBM Systems Technology Group. She has 28 years of working experience within IBM, which includes product development, L2/PFE support, and Product Engineer experience working with IBM products such as servers, networking products, IBM Network Attached Storage appliances, IBM DS6000™, IBM System Storage® N series, Storwize V7000, Storwize V7000 Unified, and IBM Flex System V7000. She also worked as technical team lead when she joined the STG worldwide N series PFE support team in Research Triangle Park, North Carolina. Eva has a System Storage Certification with IBM. She was a participant in developing IBM Storage Networking Solutions V1 and V2 Certification test. She holds a Masters degree in Computer Science.

Carsten Larsen is an IBM Certified Senior IT Specialist working for the Technical Services Support organization in IBM Denmark delivering consultancy services to IBM customers within the storage arena. Carsten joined IBM in 2007, leaving behind a job at HP where he worked with Storage Arrays and UNIX for 10 years, holding, among others, the HP certification: Master Certified Systems Engineer (MASE). While working for IBM, Carsten has obtained Brocade BCFP and BCSD certifications as well as NetApp NCDA and NCIE certifications. Carsten is the author of a number of IBM Redbooks publications related to these product certifications.

Kim Serup is an IT Specialist at STG Lab Services, Copenhagen, Denmark, and has been with IBM since 1998. Kim is delivering consultancy services to IBM customers within the PureFlex, x Series, and Storage arenas. Before Kim joined STG Lab Services, he worked as a support specialist for IBM System x PFE support and provided EMEA Level 2 support for IBM DS3x00 products. At the beginning of Kim's career within IBM, Kim worked for the Technical Services Support organization, delivering hardware support and education for IBM customers, Business partners and IBM colleagues.

John Sexton is a Certified Consulting IT Specialist, based in Auckland, New Zealand, and has over 25 years of experience working in IT. He has worked at IBM for the last 16 years. His areas of expertise include IBM System p, AIX, IBM HACMP™, virtualization, storage, IBM Tivoli® Storage Manager, SAN, SVC, and business continuity. He provides pre-sales support and technical services for clients throughout New Zealand, including consulting, solution design and implementation, troubleshooting, performance monitoring, system migration, and training. Prior to joining IBM in New Zealand, John worked in the United Kingdom supporting and maintaining systems in the financial and advertising industries.

Mansoor Syed is a Storage Support Specialist for the IBM Entry and Midrange Storage Systems, in Research Triangle Park, North Carolina. He provides hardware and software product support for IBM System Storage midrange and high performance computing systems as well as IBM virtualized disk system solutions. Mansoor joined IBM as a Client Advocate in 2004 and since then has contributed in various technical support teams for products ranging from IBM BladeCenter®, IBM xSeries®, and System x servers, to High Availability and High Performance Cluster Solutions. His responsibilities also involve enhancing RAS packaging of several IBM products, training of support engineers, and management of hardware infrastructure for Support teams. Presently, Mansoor provides technical leadership in the Product Field Engineering team for IBM System Storage DS5000 and DS3000 series.

Alexander Watson is an ATS Specialist for Storage Advanced Technical Skills (ATS) Americas in the United States. He is a Subject Matter Expert on the IBM Midrange system storage products and SAN switches. He has over fifteen years of experience in planning, managing, designing, implementing, problem analysis, and tuning of storage systems and SAN environments. He has worked at IBM for twelve years. His areas of expertise include IBM Midrange Storage solutions, SAN fabric networking, and Open System Storage I/O performance. Al has been a co-author of more than thirteen IBM Redbooks publications and IBM Redpapers publications™.

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

Jon Tate
Tamikia Barrow

International Technical Support Organization, Raleigh Center

Matt Riddle
Mary Connell
Doris Konieczny
Roger Bullard
Walter Tita
Andy McNeil
Muhammad Zubair
Andy Sylivant
John Fasano
Karen Brown
Bill Wiegand
Andrew Martin
Dan Braden

IBM

Now you can become a published author, too!

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

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 Redbooks form found at:

ibm.com/redbooks

- ▶ Send your comments in an email to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

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

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:
<http://www.facebook.com/IBMRedbooks>
- ▶ Follow us on Twitter:
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.html>



Introduction to IBM Flex Systems and IBM PureSystems offerings

This chapter provides an overview of the IBM PureSystems offerings and how the IBM Flex System V7000 Storage Node adds to a cloud ready solution within a single IBM Flex System Enterprise Chassis. Such a solution consists of Compute Nodes, Storage Systems, LAN-infrastructure, and SAN-infrastructure that connects it all together.

IBM Flex System products are ideally suited for data center environments that require flexible, cost-effective, secure, and energy-efficient hardware.

The innovative design features of the IBM Flex System products make it possible for you to configure totally integrated, customized, secure solutions that meet your data center needs today and provide flexible expansion capabilities for the future. The scalable hardware features and the unprecedented cooling capabilities of the IBM Flex System products help you optimize hardware utilization, minimize cost, and simplify the overall management of your data center.

The primary focus of this book is to describe features and functions of IBM Flex System V7000 Storage Node. However, in early versions of the IBM Flex System, the integrated storage is provided by IBM Storwize V7000. Hence this introduction covers both storage systems. In the following chapters, mainly the IBM Flex System V7000 Storage Node are described.

For more information about the IBM PureSystems, see the following website:

<http://www.ibm.com/ibm/puresystems/us/en/index.html>

1.1 IBM PureSystems overview

During the last 100 years, information technology has moved from a specialized tool to a pervasive influence on nearly every aspect of life. From tabulating machines that simply counted with mechanical switches or vacuum tubes to the first programmable computers, IBM has been a part of this growth, while always helping customers to solve problems.

Information Technology (IT) is a constant part of business and of our lives. IBM expertise in delivering IT solutions has helped the planet become smarter. And as organizational leaders seek to extract more real value from their data, business processes and other key investments, IT is moving to the strategic center of business.

To meet those business demands, IBM is introducing a new category of systems, systems that combine the flexibility of general-purpose systems, the elasticity of cloud computing and the simplicity of an appliance that is tuned to the workload. Expert integrated systems are essentially the building blocks of capability. This new category of systems represents the collective knowledge of thousands of deployments, established best practices, innovative thinking, IT leadership and distilled expertise.

The offerings in IBM PureSystems are designed to deliver value in the following ways:

- ▶ Built-in expertise helps you to address complex business and operational tasks automatically.
- ▶ Integration by design helps you to tune systems for optimal performance and efficiency.
- ▶ Simplified experience, from design to purchase to maintenance, creates efficiencies quickly.

The IBM PureSystems offerings are optimized for performance and virtualized for efficiency. These systems offer a no-compromise design with system-level upgradeability. IBM PureSystems is built for cloud, containing “built-in” flexibility and simplicity.

At IBM, expert integrated systems come in two types:

- ▶ IBM PureFlex System: Infrastructure systems deeply integrate the IT elements and expertise of your system infrastructure.
- ▶ IBM PureApplication™ System: Platform systems include middleware and expertise for deploying and managing your application platforms.

IBM PureSystems are built for cloud with integrated elasticity and virtualization capabilities to provision new services in minutes and improve business flexibility while reducing cost.

IBM Flex System is a build-to-order offering that is integrated by the client or a partner and does not deliver against all of the three attributes of expert integrated systems (built-in expertise, integration by design, simplified experience). IBM Flex System allows clients to build their own system to meet unique IT requirements with a set of no-compromise components including compute, storage, networking, and systems management.

IBM PureFlex System and IBM PureApplication System are built on elements of the IBM Flex System. It has been designed for clients that need pre-integrated hardware infrastructure comprised of compute, storage and networking nodes as well as a choice of operating systems and hypervisors.

The new IBM Flex System V7000 Storage Node is supported with IBM PureFlex System and other IBM Flex System configurations.

1.1.1 Product names

The primary product names for the IBM PureSystems components are as follows:

- ▶ IBM PureSystems:
 - The overall name for IBM's new family of expert integrated systems
- ▶ IBM Flex System:
 - A build-to-order offering with clients' choice of IBM Flex System components
 - IBM Flex System that can help you go beyond blades
 - An innovative Enterprise Chassis designed for new levels of simplicity, flexibility, integration, reliability, and upgradability
 - A broad range of x86 and IBM POWER® compute nodes
 - New IBM Flex System V7000 Storage Node built into the Enterprise Chassis
- ▶ IBM PureFlex System:
 - A solution that combines compute nodes, storage, networking, virtualization and management into a single infrastructure system. It is expert at sensing and anticipating resource needs to optimize your infrastructure.
- ▶ IBM PureApplication System:
 - A platform system designed and tuned specifically for transactional web and database applications. Its workload-aware, flexible platform is designed to be easy to deploy, customize, safeguard and manage.
- ▶ IBM Flex System V7000 Storage Node:
 - The product name for the IBM Flex System V7000 Storage Node family of controller and expansion enclosures. The IBM Flex System V7000 Storage Node is an add-on for the IBM Flex System Enterprise Chassis.
- ▶ IBM Flex System V7000 Control Enclosure:
 - The controller enclosure of the IBM Flex System V7000 Storage Node. The IBM Flex System V7000 Control Enclosure is an add-on for the IBM Flex System Enterprise Chassis and mounts internally into it.
 - The IBM Flex System V7000 Control Enclosure provides 24 disk drive bays.
 - The IBM Flex System V7000 Control Enclosure supports block workloads only.
- ▶ IBM Flex System V7000 Expansion Enclosure:
 - A SAS disk shelf with 24 disk drive bays that connects to the IBM Flex System V7000 Control Enclosure. The IBM Flex System V7000 Expansion Enclosure is an add-on for the IBM Flex System Enterprise Chassis and mounts internally into it.
- ▶ IBM Storwize V7000:
 - The IBM Storwize V7000 is a disk system with built-in IBM SAN Volume Controller (SVC) functionality. It has the ability to virtualize a wide range of external storage systems from either IBM or other Storage vendors.
 - The IBM Storwize V7000 Control Enclosure provides a choice of 24 x 2.5" Small Form Factor (SFF) disk drives or 12 x 3.5" Large Form Factor (LFF) disk drive form factors.
 - The IBM Storwize V7000 supports block workloads only.
- ▶ IBM Storwize V7000 Unified:
 - IBM Storwize V7000 Unified is like the IBM Storwize V7000 a disk system that provides internal storage and external virtualization. However, the IBM Storwize V7000 Unified also has file modules that provide NAS functionality like the CIFS and NFS protocols.
 - The Storwize V7000 Unified consolidates block and file workloads into a single system.

- ▶ IBM Storwize V7000 Control Enclosure:
 - This component is the controller enclosure of the IBM Storwize V7000 storage system.
 - The IBM Storwize V7000 Control Enclosure provides 12 or 24 disk drive bays, depending on the model.
- ▶ IBM Storwize V7000 Expansion Enclosure:
 - A SAS disk shelf with either 12 or 24 disk drive bays that can connect to either the IBM Storwize V7000 Control Enclosure or the IBM Flex System V7000 Control Enclosures.

Figure 1-2 shows the different IBM PureSystems and their building blocks.




Figure 1-1 IBM PureSystems

1.1.2 IBM PureFlex System

To meet today’s complex and ever-changing business demands, you need a solid foundation of server, storage, networking, and software resources that is simple to deploy and can quickly and automatically adapt to changing conditions. You also need access to, and the ability to take advantage of, broad expertise and proven best practices in systems management, applications, hardware maintenance, and more.

IBM PureFlex System is a comprehensive infrastructure system that provides an expert integrated computing system, combining servers, enterprise storage, networking, virtualization, and management into a single structure. Its built-in expertise enables organizations to simply manage and flexibly deploy integrated patterns of virtual and hardware resources through unified management. These systems are ideally suited for customers interested in a system that delivers the simplicity of an integrated solution but who also want control over tuning middleware and the run-time environment.

IBM PureFlex System recommends workload placement based on virtual machine compatibility and resource availability. Using built-in virtualization across servers, storage, and networking, the infrastructure system enables automated scaling of resources and true workload mobility.

IBM PureFlex System undergoes significant testing and experimentation, so it can mitigate IT complexity without compromising the flexibility to tune systems to the tasks businesses demand. By providing both flexibility and simplicity, IBM PureFlex System can provide extraordinary levels of IT control, efficiency, and operating agility that enable businesses to rapidly deploy IT services at a reduced cost. Moreover, the system is built on decades of expertise, enabling deep integration and central management of the comprehensive, open-choice infrastructure system and dramatically cutting down on the skills and training required for managing and deploying the system.

IBM PureFlex System combine advanced IBM hardware and software along with patterns of expertise and integrate them into three optimized configurations that are simple to acquire and deploy so you get fast time to value for your solution.

Figure 1-2 shows the IBM PureFlex System with its three different chassis implementations.

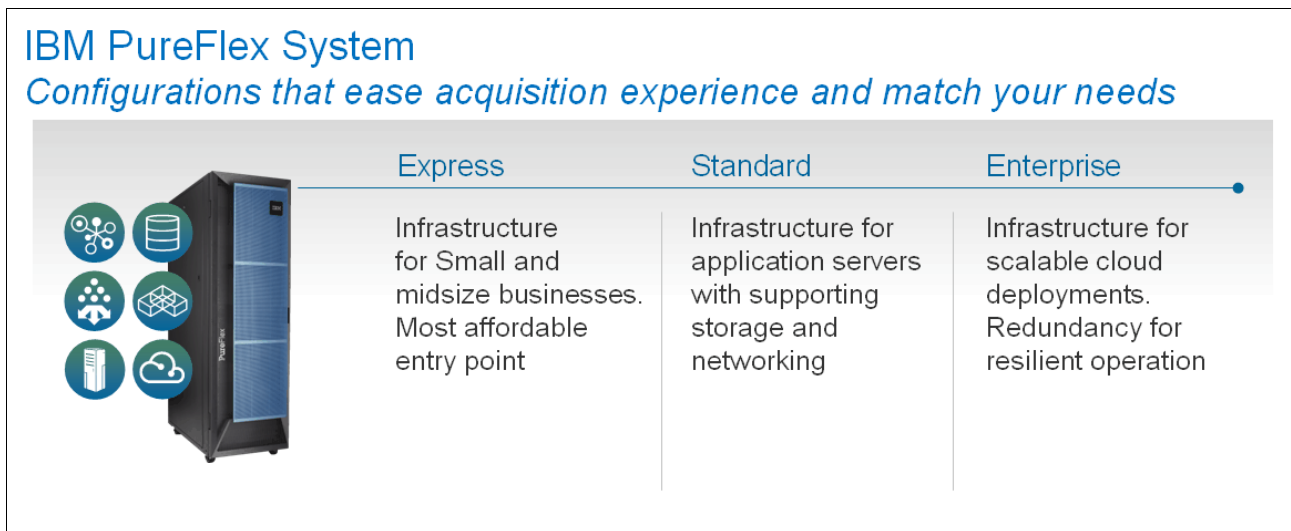


Figure 1-2 IBM PureFlex System

The three PureFlex System configurations are as follows:

- ▶ IBM PureFlex System Express: Designed for small and medium businesses, it is the most affordable entry point for the PureFlex System.
- ▶ IBM PureFlex System Standard: Optimized for application servers with supporting storage and networking, it is designed to support your key ISV solutions.
- ▶ IBM PureFlex System Enterprise: Optimized for transactional and database systems, it has built-in redundancy for highly reliable and resilient operation to support your most critical workloads.

Tip: IBM Flex System allows you to build your own system to meet the unique IT requirements.

The components of the PureFlex System are summarized in Table 1-1.

Table 1-1 IBM PureFlex System components

Component	IBM PureFlex System Express	IBM PureFlex System Standard	IBM PureFlex System Enterprise
IBM PureFlex System 42U Rack	1	1	1
IBM Flex System Enterprise Chassis	1	1	1
IBM Flex System Fabric EN4093 10Gb Scalable Switch	1	1	2 with both port-count upgrades
IBM Flex System FC3171 8 Gb SAN Switch	1	2	2
IBM Flex System Manager Node	1	1	1
IBM Flex System Manager software license	IBM Flex System Manager with 1-year service and support	IBM Flex System Manager Advanced with 3-year service and support	IBM Flex System Manager Advanced with 3-year service and support
IBM Flex System Chassis Management Module	2	2	2
Chassis power supplies (std/max)	2 / 6	4 / 6	6 / 6
IBM Flex System Enterprise Chassis 80 mm Fan Module Pair (std/max)	4 / 8	6 / 8	8 / 8
IBM Flex System V7000 Storage Node	Yes (redundant controller)	Yes (redundant controller)	Yes (redundant controller)
IBM Flex System V7000 Base Software	Base with 1-year software maintenance agreement	Base with 3-year software maintenance agreement	Base with 3-year software maintenance agreement

The fundamental building blocks of IBM PureFlex System solutions are the IBM Flex System Enterprise Chassis complete with compute nodes, networking, and storage. See the next sections for more information about the building blocks of the IBM PureFlex System.

- ▶ 1.2, “IBM PureFlex System building blocks” on page 8
- ▶ 1.3, “IBM Flex System Enterprise Chassis” on page 10
- ▶ 1.4, “Compute nodes” on page 15
- ▶ 1.5, “I/O modules” on page 24

1.1.3 IBM PureApplication System

The IBM PureApplication System is a platform system that pre-integrates a full application platform set of middleware and expertise in with the IBM PureFlex System with a single management console. It is a workload-aware, flexible platform that is designed to be easy to deploy, customize, safeguard, and manage in a traditional or private cloud environment, ultimately providing superior IT economics.

Availability: IBM Flex System V7000 Storage Node is currently not offered in IBM PureApplication Systems. Currently, the only available storage for IBM PureApplication Systems is IBM Storwize V7000.

With the IBM PureApplication System, you can provision your own patterns of software, middleware, and virtual system resources. You can provision these patterns within a unique framework that is shaped by IT best practices and industry standards. Such standards have been developed from many years of IBM experience with clients and from a deep understanding of smarter computing. These IT best practices and standards are infused throughout the system.

With IBM PureApplication System, you enjoy the following benefits:

- ▶ IBM builds expertise into preintegrated deployment patterns, which can speed the development and delivery of new services.
- ▶ By automating key processes such as application deployment, PureApplication System built-in expertise capabilities can reduce the cost and time required to manage an infrastructure.
- ▶ Built-in application optimization expertise reduces the number of unplanned outages through best practices and automation of the manual processes identified as sources of those outages.
- ▶ Administrators can use built-in application elasticity to scale up or to scale down automatically. Systems can use data replication to increase availability.

Patterns of expertise can automatically balance, manage, and optimize the elements necessary, from the underlying hardware resources up through the middleware and software. These patterns of expertise help deliver and manage business processes, services, and applications by encapsulating best practices and expertise into a repeatable and deployable form. This best-practice knowledge and expertise has been gained from decades of optimizing the deployment and management of data centers, software infrastructures, and applications around the world.

These patterns help you achieve the following types of value:

- ▶ **Agility:** As you seek to innovate to bring products and services to market faster, you need fast time-to-value. Expertise built into a solution can eliminate manual steps, automate delivery, and support innovation.
- ▶ **Efficiency:** To reduce costs and conserve valuable resources, you must get the most out of your systems with energy efficiency, simple management, and fast, automated response to problems. With built-in expertise, you can optimize your critical business applications and get the most out of your investments.
- ▶ **Increased simplicity:** You need a less complex environment. Patterns of expertise help you to easily consolidate diverse servers, storage and applications onto an easier-to-manage, integrated system.
- ▶ **Control.** With optimized patterns of expertise, you can accelerate cloud implementations to lower risk by improving security and reducing human error.

IBM PureApplication System is available in four configurations. These configuration options enable you to choose the size and compute power that meets your needs for application infrastructure. You can upgrade to the next size when your organization requires more capacity, and in most cases, you can do so without application downtime.

Table 1-2 provides a high-level overview of the configurations.

Table 1-2 IBM PureApplication System configurations

	IBM PureApplication System W1500-96	IBM PureApplication System W1500-192	IBM PureApplication System W1500-384	IBM PureApplication System W1500-608
Cores	96	192	384	608
RAM	1.5 TB	3.1 TB	6.1 TB	9.7 TB
SSD Storage	6.4 TB			
HDD Storage	48.0 TB			
Application Services Entitlement	Included			

For more details about IBM PureApplication System, see the following website:

<http://ibm.com/expert>

1.2 IBM PureFlex System building blocks

IBM PureFlex System provides an integrated computing system, combining servers, enterprise storage, networking, virtualization, and management into a single structure. The built-in expertise lets organizations simply manage and flexibly deploy integrated patterns of virtual and hardware resources through unified management.

1.2.1 Highlights

Each system consists of System x nodes, IBM Power Systems™ compute nodes, or a combination of these two types, which is known as a hybrid configuration. The bundled, on-site services provide some initial compute node configuration and might differ for System x nodes and Power Systems compute nodes. A client-specified primary node (POWER or x86) is pre-configured with a hypervisor (IBM PowerVM®, VMWare, KVM, HyperV) to allow virtual server configuration by IBM services personnel. Services also include skills transfer to the client personnel.

Important: Initial IBM PureFlex System configuration is carried out by IBM services and is included with the purchase. To ensure configuration success, the default shipped configuration must not be changed until these services are completed by IBM.

1.2.2 Components

The IBM PureFlex System offerings comprise of the following components as illustrated in Figure 1-3. With these components pre-configured, pre-integrated infrastructure systems with compute, storage, networking, physical and virtual management, and entry cloud management with integrated expertise are delivered by the IBM PureFlex System.




Figure 1-3 IBM PureFlex System Building Blocks

Storage components

The storage capabilities of IBM Flex System allows you to gain advanced functionality with storage nodes in your system while taking advantage of your existing storage infrastructure through advanced virtualization. For early versions of the IBM Flex System, the only integrated storage was the IBM Storwize V7000 that was external to the IBM Flex System Enterprise Chassis. With the introduction of the IBM Flex System V7000 Storage Node, storage is provided internally from the IBM Flex System Enterprise Chassis.

Simplified management

The IBM Flex System simplifies storage administration with a single user interface for all your storage with a management console that is integrated with the comprehensive management system. These management and storage capabilities allow you to virtualize third-party storage with non-disruptive migration of the current storage infrastructure. You can also take advantage of intelligent tiering so you can balance performance and cost for your storage needs. The solution also supports local and remote replication and snapshots for flexible business continuity and disaster recovery capabilities.

Infrastructure

The IBM Flex System Enterprise Chassis is the foundation of the offering, supporting intelligent workload deployment and management for maximum business agility. The 10U high chassis has the capacity of up to 14 compute nodes, or a mix of compute nodes and Storage, which mounts from the front. From the rear, it mounts power supplies, fans, and different options of LAN and SAN switches. The IBM Flex System Enterprise Chassis delivers high-performance connectivity for your integrated compute nodes, storage, networking, and management resources. The chassis is designed to support multiple generations of technology and offers independently scalable resource pools for higher utilization and lower cost per workload.

We now review the various components of the IBM Flex System in order to understand how the IBM Flex System V7000 Storage Node integrates with the PureFlex Systems solution. All of the components are used in the three pre-integrated offerings to support compute, storage, and networking requirements. You can select from these offerings, which are designed for key client initiatives and help simplify ordering and configuration.

While we only provide a summary of the IBM Flex System components in the following sections, its important to understand the various available options for IBM Flex System before we describe the IBM Flex System V7000 Storage Node in detail in Chapter 2, “Introduction to IBM Flex System V7000 Storage Node” on page 39.

For detailed information about the components, see *IBM PureFlex System and IBM Flex System Products and Technology*, SG24-7984.

1.3 IBM Flex System Enterprise Chassis

The IBM Flex System Enterprise Chassis is a 10U next-generation server platform with integrated chassis management. It is a compact, high-density, high-performance, rack-mount, scalable server platform system. It supports up to 14 one-bay compute nodes that can share common resources, such as power, cooling, management, and I/O resources within a single Enterprise Chassis. In addition, it can also support up to seven 2-bay compute nodes or three 4-bay compute nodes (three IBM Flex System V7000 Storage Nodes or expansion enclosures) when the shelves are removed from the chassis. The 1-bay, 2-bay, and 4-bay components can be “mixed and matched” to meet specific hardware requirements.

Figure 1-4 shows the chassis with the IBM Flex System V7000 Storage Node occupying 4 x compute bays, which is partially inserted into the chassis for identification.



Figure 1-4 Front view of IBM Enterprise Flex System Chassis with an IBM Flex System V7000 Storage Node

The chassis has the following features on the front:

- ▶ The front information panel, located on the lower left of the chassis
- ▶ Bays 1 to 14 supporting Nodes, storage enclosures and FSM
- ▶ Lower airflow inlet apertures that provides air cooling for switches, IBM Flex System Chassis Management Module and power supplies
- ▶ Upper airflow inlet apertures that provide cooling for power supplies

For proper cooling, each bay in the front or rear of the chassis must contain either a device or a filler.

The Enterprise Chassis provides several LEDs on the front information panel that can be used to obtain the status of the chassis. The Identify, Check log and the Fault LED also appear on the rear of the chassis for ease of use.

The major components of Enterprise Chassis are as follows:

- ▶ Fourteen 1-bay compute node bays (can also support seven 2-bay or three 4-bay compute nodes with shelves removed).
- ▶ Six 2500-watt power modules that provide N+N or N+1 redundant power.

- ▶ Ten fan modules (eight 80 mm fan modules and two 40 mm fan modules).
- ▶ Four physical I/O modules.
- ▶ An I/O architectural design capable of providing:
 - Up to 8 lanes of I/O to an I/O adapter card; each lane capable of up to 16 Gbps
 - A maximum of 16 lanes of I/O to a half wide-node with two adapters
 - A wide variety of networking solutions including Ethernet, Fibre Channel, FCoE, and InfiniBand
- ▶ Up to two IBM Flex System Manager (FSM) management appliances for redundancy. The FSM provides multiple-chassis management support for up to four chassis.
- ▶ Two IBM Flex System Chassis Management Module (CMMs). The CMM provides single-chassis management support.

The chassis can be configured with the following information about the chassis location:

- ▶ Rack Room
- ▶ Rack Location
- ▶ Position in Rack (the lowest Rack Unit occupied by the Chassis)
- ▶ Chassis Name
- ▶ Bay ID

Individual components will then be able to work out their bay in the chassis, the IBM Flex System V7000 Storage Node enclosure uses 4 bays (double wide and double tall) per enclosure and will report its bay as the lowest left bay that it occupies.

Figure 1-5 shows the rear of the chassis where the I/O modules and chassis management modules can be seen.




Figure 1-5 Rear view of the IBM Enterprise Flex System Chassis

The following components can be installed into the rear of the chassis:

- ▶ Up to two IBM Flex System Chassis Management Modules (CMM).
- ▶ Up to six 2500W power supply modules.
- ▶ Up to six fan modules consisting of four 80 mm fan modules and two 40 mm fan modules. Additional fan modules can be installed, up to a total of ten modules.
- ▶ Up to four I/O modules.

1.3.1 Chassis power supplies

A maximum of six power supplies can be installed within the Enterprise Chassis. The PSUs and empty PSU bays can be seen in Figure 1-5 on page 12. The power supplies are 80 PLUS Platinum certified and are 2100 Watts and 2500 Watts output rated at 200VAC, with oversubscription to 2895 Watts and 3538 Watts output at 200VAC respectively. The power supply operating range is 200-240VAC. The power supplies also contain two independently powered 40 mm cooling fan modules that pick power up from the midplane, not from the power supply.

Availability: The 2100W power supplies are only available by Configure to Order (CTO). For more information about the 2100W power supply, see *IBM PureFlex System and IBM Flex System Products and Technology*, SG24-7984.

Highlights

The chassis allows configurations of power supplies to give N+N or N+1 redundancy. A fully configured chassis will operate on just three 2500 W power supplies with no redundancy, but N+1 or N+N is likely to be preferred. Using three (or six with N+N redundancy) power supplies allows for a balanced 3-phase configuration.

All power supply modules are combined into a single power domain within the chassis, which distributes power to each of the compute nodes, I/O modules, and ancillary components through the Enterprise Chassis midplane. The midplane is a highly reliable design with no active components. Each power supply is designed to provide fault isolation and is hot swappable.

There is power monitoring of both the DC and AC signals from the power supplies, which allows the IBM Flex System Chassis Management Module to accurately monitor these signals. The integral power supply fans are not dependent upon the power supply being functional, they operate and are powered independently from the midplane.

Each power supply in the chassis has a 16A C20 3 pin socket and can be fed by a C19 power cable, from a suitable supply.

The chassis power system is designed for efficiency using datacenter power consisting of 3 phase 60A Delta 200 VAC (North America) or 3 phase 32A wye 380-415 VAC (international). The Chassis can also be fed from single phase 200-240 VAC supplies if required.

Power supply redundancy

There are different vendors that can vary slightly in the terminology when describing power supply unit (PSU) redundancy. In general, ' N ' is the minimum number of PSUs required to keep the server operational, in this case, to keep the populated IBM Flex System Enterprise Chassis operational. The term ' $N+1$ ' is the minimum number of PSUs plus one. It is not the best option for redundancy, it is the equivalent of a "hot spare" and will protect against PSU failure only.

The minimum number of PSUs required to keep the server or chassis operational duplicated or doubled is referred to as ' $N+N$ '. The PSUs are fully redundant where there will be an even number for minimum ' $N+N$ ' support.

An IBM Flex System Enterprise Chassis would typically be connected to at least two power distribution units (PDU) in a computer room with the PSUs connected evenly across the PDUs. With ' $N+N$ ' redundancy the IBM Flex System Enterprise Chassis in this configuration is also protected against PDU failure within the computer room.

1.3.2 Fan modules and cooling

The Enterprise Chassis supports up to ten hot pluggable fan modules consisting of two 40 mm fan modules and eight 80 mm fan modules.

Highlights

A chassis can operate with a minimum of six hot-swap fan modules installed, consisting of four 80 mm fan modules and two 40 mm fan modules. The fan modules plug into the chassis and connect to the fan distribution cards. The 80 mm fan modules can be added as required to support chassis cooling requirements.

The two 40 mm fan modules in fan bays 5 and 10 (top two) distribute airflow to the I/O modules and chassis management modules. These modules ship pre installed in the chassis.

Each 40 mm fan module contains two 40 mm fans internally, side by side.

The 80 mm fan modules distribute airflow to the compute nodes through the chassis from front to rear. Each 80 mm fan module contains two 80 mm fan modules, back to back at each end of the module, which are counter rotating.

Both fan modules have an EMC (electromagnetic compatibility) mesh screen on the rear internal face of the module. The design of this also has an additional benefit for the airflow, by providing a laminar flow through the screen, which reduces turbulence of the exhaust air and improves the efficiency of the overall fan assembly. Laminar flow is a smooth flow of air, sometimes called streamline flow. The opposite of a laminar flow is a turbulent flow. The design of the whole fan assembly, the fan blade design, the distance between and size of the fan modules together with the EMC mesh screen ensures a highly efficient fan design that provides the best cooling for lowest energy input.

The minimum number of 80 mm fan modules is four. The maximum number of 80 mm fan modules that can be installed is eight. When the modules are ordered as an option, they are supplied as a pair.

Environmental specifications

The chassis is designed to operate in temperatures up to 40°C (104°F), in ASHRAE class A3 operating environments.

The airflow requirements for the Enterprise Chassis are from 270 CFM (cubic feet per minute) to a maximum of 1020 CFM.

Environmental specifications are as follows:

- ▶ Humidity, non-condensing: -12°C dew point (10.4°F) and 8% - 85% relative humidity
- ▶ Maximum dew point: 24°C (75°F)
- ▶ Maximum elevation: 3050 m (10,006 ft.)
- ▶ Maximum rate of temperature change: 5°C/hr. (41°F/hr.)

Heat output (approximate):

- ▶ Maximum configuration: potentially 12.9kW

The 12.9 kW figure is only a potential maximum, where the most power hungry configuration is chosen and all power envelopes are maximum. For a more realistic figure, the IBM Power Configurator tool can be made to establish specific power requirements for a given configuration.

The Power Configurator tool can be found at the following website:

<http://www.ibm.com/systems/x/hardware/configtools.html>

1.4 Compute nodes

The IBM Flex System portfolio of compute nodes includes those with Intel Xeon processors or with IBM POWER7® processors. Depending on the compute node design, it can come in one of two different form factors:

- ▶ Half-wide node: Occupies one chassis bay, half the width of the chassis (approximately 215 mm or 8.5").
- ▶ Full-wide node: Occupies two chassis bays side-by-side, the full width of the chassis (approximately 435 mm or 17").

The applications installed on the compute nodes can be running on an operating system run natively on a dedicated physical server or can be virtualized in a virtual machine managed by a hypervisor layer. Here we provide a summary of the compute nodes. For further detailed information about these topics, see the *IBM Flex System p260 and p460 Planning and Implementation Guide*, SG24-7989 and *IBM PureFlex System and IBM Flex System Products and Technology*, SG24-7984.

1.4.1 IBM Flex System x440 Compute Node

The IBM Flex System x440 Compute Node (machine type 7917) is a high-density four socket server, optimized for high-end virtualization, mainstream database deployments, memory-intensive and high performance environments.

The IBM Flex System x440 Compute Node is a double-wide compute node providing scalability to support up to four Intel Xeon E5-4600 processors. The node's width allows for a significant I/O capability.

Figure 1-6 shows the front of the compute node, showing the location of the controls, LEDs, and connectors. The light path diagnostic panel is located on the upper edge of the front panel bezel, in the same place as the x220 and x240.




Figure 1-6 IBM Flex System x440 Compute Node

See *IBM Flex System x440 Compute Node*, TIPS0886 for more information.

1.4.2 IBM Flex System x240 Compute Node

The IBM Flex System x240 Compute Node, available as machine type 8737, is a half-wide, two-socket server running the latest Intel Xeon processor E5-2600 family processors. It is ideal for infrastructure, virtualization, and enterprise business applications and is compatible with the IBM Flex System Enterprise Chassis. The x240 supports up to two Intel Xeon E5-2600 series multi-core processors, 24 DIMM modules, two hot-swap drives, two PCI Express I/O adapter cards, and has an option for two internal USB connectors. Figure 1-7 shows the single bay x240 compute node.



Figure 1-7 The x240 type 8737

The IBM Flex System x240 Compute Node type 8737 features the Intel Xeon E5-2600 series processors with two, four, six, or eight cores per processor, with up to 16 threads per socket. The processors have up to 20 MB of shared L3 cache, Hyper-Threading, Turbo Boost Technology 2.0 (depending on processor model), two QuickPath Interconnect (QPI) links that run at up to 8 GT/s, one integrated memory controller, and four memory channels supporting up to three DIMMs each.

The x240 includes 8 GB of memory (2 x 4 GB DIMMs) running at either 1600 MHz or 1333 MHz depending on model. Some models include an Embedded 10 Gb Virtual Fabric Ethernet LAN-on-motherboard (LOM) controller as standard; this embedded controller precludes the use of an I/O adapter in I/O connector 1. Model numbers in the form x2x (for example, 8737-L2x) include an Embedded 10 Gb Virtual Fabric Ethernet LAN-on-motherboard (LOM) controller as standard. Model numbers in the form x1x (for example, 8737-A1x) do not include this embedded controller.

The x240 with the Intel Xeon E5-2600 series processors can support up to 768 GB of memory in total when using 32 GB LRDIMMs and with both processors installed. The x240 uses Double Data Rate-3 (DDR-3) low-profile (LP) DIMMs. The x240 supports three types of DIMM memory:

- ▶ Registered DIMM (RDIMM) modules
- ▶ Unbuffered DIMM (UDIMM) modules
- ▶ Load-reduced (LRDIMM) modules

The mixing of these different memory DIMM types is not supported.

The x240 compute node features an onboard LSI 2004 SAS controller with two small form factor (SFF) hot-swap drive bays that are accessible from the front of the compute node. The onboard LSI SAS2004 controller provides RAID 0, RAID 1, or RAID 10 capability and supports up to two SFF hot-swap SAS or SATA HDDs or two SFF hot-swap solid state drives. Figure 1-8 shows how the LSI2004 SAS controller and hot-swap storage devices connect to the internal HDD interface.




Figure 1-8 The LSI2004 SAS controller connections to HDD interface

Each x240 server has an Integrated Management Module version 2 (IMMv2) onboard and uses the Unified Extensible Firmware Interface (UEFI) to replace the older BIOS interface.

Embedded 10 Gb Virtual Fabric Adapter

Some models of the x240 include an Embedded 10 Gb Virtual Fabric Adapter (VFA, also known as LAN on Motherboard or LOM), built into the system board. Each of these models that includes the embedded 10 Gb VFA also has the Compute Node Fabric Connector installed in I/O connector 1 (and physically screwed onto the system board) to provide connectivity to the Enterprise Chassis midplane.

I/O expansion

The x240 has two PCIe 3.0 x16 I/O expansion connectors for attaching I/O adapter cards. There is also another expansion connector designed for future expansion options. The I/O expansion connectors are a very high-density 216 pin PCIe connector. By installing I/O adapter cards, it allows the x240 to connect with switch modules in the IBM Flex System Enterprise Chassis.

Figure 1-9 shows the rear of the x240 compute node and the locations of the I/O connectors.




Figure 1-9 Rear of the x240 compute node showing the locations of the I/O connectors

1.4.3 IBM Flex System x220 Compute Node

The IBM Flex System x220 Compute Node, machine type 7906, is the next generation cost-optimized compute node designed for less demanding workloads and low-density virtualization. The x220 is efficient and equipped with flexible configuration options and advanced management to run a broad range of workloads. It is a high-availability, scalable compute node optimized to support the next-generation microprocessor technology. With a balance between cost and system features, the x220 is an ideal platform for general business workloads. This section describes the key features of the server.

Highlights

Figure 1-10 shows the front of the compute node indicating location of the controls, LEDs, and connectors.




Figure 1-10 IBM Flex System x220 Compute Node

The IBM Flex System x220 Compute Node features the Intel Xeon E5-2400 series processors. The Xeon E5-2400 series processor has models with either four, six, or eight cores per processor with up to 16 threads per socket. The processors have up to 20 MB of shared L3 cache, Hyper-Threading, Turbo Boost Technology 2.0 (depending on processor model), one QuickPath Interconnect (QPI) link that runs at up to 8 GT/s, one integrated memory controller, and three memory channels supporting up to two DIMMs each.

The x220 also supports an Intel Pentium 1403 or 1407 dual-core processor for entry-level server applications. Only one Pentium processor is supported in the x220. CPU socket 2 must be left unused and only six DIMM socks are available.

The x220 supports Low Profile (LP) DDR3 memory registered DIMMs (RDIMMs) and unbuffered DIMMs (UDIMMs). The server supports up to six DIMMs when one processor is installed and up to 12 DIMMs when two processors are installed. Each processor has three memory channels, and there are two DIMMs per channel.

The x220 server has two 2.5-inch hot-swap drive bays accessible from the front of the blade server as shown in Figure 1-10. The server optionally supports three internal disk controllers allowing a greater number of internal drives up to a maximum of eight with the ServerRAID M5115 controller and also supports 1.8-inch solid-state drives.

Each IBM Flex System x220 Compute Node has an Integrated Management Module version 2 (IMMv2) onboard and uses the Unified Extensible Firmware Interface (UEFI).

Embedded 1 Gb Ethernet controller

Some models of the x220 include an Embedded 1 Gb Ethernet controller (also known as LAN on Motherboard or LOM) built into the system board. Each x220 model that includes the controller also has the Compute Node Fabric Connector installed in I/O connector 1 (and physically screwed onto the system board) to provide connectivity to the Enterprise Chassis midplane.

The Fabric Connector enables port 1 on the controller to be routed to I/O module bay 1 and port 2 to be routed to I/O module bay 2. The Fabric Connector can be unscrewed and removed, if required, to allow the installation of an I/O adapter on I/O connector 1.

1.4.4 IBM Flex System p260 and p24L Compute Nodes

The IBM Flex System p260 Compute Node and IBM Flex System p24L Compute Node are based on IBM POWER architecture technologies. These compute nodes run in IBM Flex System Enterprise Chassis units to provide a high-density, high-performance compute node environment, using advanced processing technology. The IBM Flex System p24L Compute Node shares several similarities to the IBM Flex System p260 Compute Node in that it is a half-wide, Power Systems compute node with two POWER7 processor sockets, 16 memory slots, two I/O adapter slots, and an option for up to two internal drives for local storage. The IBM Flex System p24L Compute Node is optimized for lower-cost Linux installations.

Highlights

The IBM Flex System p260 Compute Node has the following features:

- ▶ Two processors with up to 16 POWER7 processing cores, up to 8 per processor
- ▶ Sixteen DDR3 memory DIMM slots supporting IBM Active Memory™ Expansion
- ▶ Supports VLP (Very Low Profile) and LP (Low Profile) DIMMs
- ▶ Two P7IOC I/O hubs
- ▶ RAID-compatible SAS controller supporting up to 2 SSD or HDD drives
- ▶ Two I/O adapter slots
- ▶ Flexible Support Processor (FSP)
- ▶ System management alerts
- ▶ IBM Light Path Diagnostics
- ▶ USB 2.0 port
- ▶ IBM EnergyScale™ technology

The front panel of Power Systems compute nodes has the following common elements, as shown in Figure 1-11:

- ▶ USB 2.0 port
- ▶ Power-control button and light path, light-emitting diode (LED) (green)
- ▶ Location LED (blue)
- ▶ Information LED (amber)
- ▶ Fault LED (amber)




Figure 1-11 IBM Flex System p260 Compute Node with front panel details

There is no onboard video capability in the Power Systems compute nodes. The machines have been designed to be accessed using Serial Over LAN (SOL) or the IBM Flex System Manager (FSM).

The IBM POWER7 processor represents a leap forward in technology and associated computing capability. The multi-core architecture of the POWER7 processor has been matched with a wide range of related technologies to deliver leading throughput, efficiency, scalability, and reliability, availability, and serviceability (RAS).

Although the processor is an important component in servers, many elements and facilities have to be balanced across a server to deliver maximum throughput. As with previous generations of systems based on POWER processors, the design philosophy for POWER7 processor-based systems is one of system-wide balance in which the POWER7 processor plays an important role.

Each POWER7 processor has an integrated memory controller. Industry standard DDR3 Registered DIMM (RDIMM) technology is used to increase reliability, speed, and density of memory subsystems.

The p260 and p24L has an onboard SAS controller that can manage up to two, non-hot-pluggable internal drives. Both 2.5-inch hard disk drives (HDDs) and 1.8-inch solid-state drives (SSDs) are supported. The maximum number of drives that can be installed in the p260 or p24L is two. SSD and HDD drives cannot be mixed.

There are several advanced system management capabilities built into the p260 and p24L. A Flexible Service Processor handles most of the server-level system management. It has features, such as system alerts and Serial-Over-LAN capability.

A Flexible Service Processor (FSP) provides out-of-band system management capabilities, such as system control, run-time error detection, configuration, and diagnostics. Generally, you do not interact with the FSP directly but, rather, using tools, such as IBM Flex System Manager, IBM Flex System Chassis Management Module, and external IBM Systems Director Management Console. The FSP provides a Serial-over-LAN interface, which is available with the IBM Flex System Chassis Management Module and the `console` command.

The p260 and p24L do not have an on-board video chip and do not support keyboard, video, and mouse (KVM) connection. Server console access is obtained by a SOL connection only. SOL provides a means to manage servers remotely by using a command-line interface (CLI) over a Telnet or secure shell (SSH) connection. SOL is required to manage servers that do not have KVM support or that are attached to the IBM Flex System Manager. SOL provides console redirection for both System Management Services (SMS) and the server operating system. The SOL feature redirects server serial-connection data over a LAN without requiring special cabling by routing the data using the IBM Flex System Chassis Management Module network interface. The SOL connection enables Power Systems compute nodes to be managed from any remote location with network access to the IBM Flex System Chassis Management Module.

The IBM Flex System Chassis Management Module CLI provides access to the text-console command prompt on each server through a SOL connection, enabling the p260 and p24L to be managed from a remote location.

I/O adapter slots

There are two I/O adapter slots on the p260 and the p24L which are identical in shape (form factor). Also different is that the I/O adapters for the Power Systems compute nodes have their own connector that plugs into the IBM Flex System Enterprise Chassis midplane.

The I/O is controlled by two P7-IOC I/O controller hub chips. This provides additional flexibility when assigning resources within Virtual I/O Server (VIOS) to specific Virtual Machine/LPARs.

1.4.5 IBM Flex System p460 Compute Node

The IBM Flex System p460 Compute Node is also based on IBM POWER architecture technologies. This compute node is a full-wide, Power Systems compute node with four POWER7 processor sockets, 32 memory slots, four I/O adapter slots, and an option for up to two internal drives for local storage. It runs in IBM Flex System Enterprise Chassis units to provide a high-density, high-performance compute node environment, using advanced processing technology.

Highlights

The IBM Flex System p460 Compute Node has the following features:

- ▶ Four processors with up to 32 POWER7 processing cores
- ▶ Thirty-two DDR3 memory DIMM slots that support IBM Active Memory Expansion
- ▶ Supports Very Low Profile (VLP) and Low Profile (LP) DIMMs
- ▶ Four P7IOC I/O hubs
- ▶ RAID-capable SAS controller that support up to two SSD or HDD drives
- ▶ Four I/O adapter slots
- ▶ Flexible Support Processor (FSP)
- ▶ System management alerts
- ▶ IBM Light Path Diagnostics
- ▶ USB 2.0 port
- ▶ IBM EnergyScale technology

The front panel of Power Systems compute nodes has the following common elements, as shown by the p460 in Figure 1-12:

- ▶ USB 2.0 port
- ▶ Power-control button and light path, light-emitting diode (LED) (green)
- ▶ Location LED (blue)
- ▶ Information LED (amber)
- ▶ Fault LED (amber)




Figure 1-12 IBM Flex System p460 Compute Node showing front indicators

The USB port on the front of the Power Systems compute nodes is useful for a variety of tasks, including out-of-band diagnostics, hardware RAID setup, operating system access to data on removable media, and local OS installation. It might be helpful to obtain a USB optical (CD or DVD) drive for these purposes, in case the need arises, as there is no optical drive in the IBM Flex System Enterprise Chassis.

Although the processor is an important component in servers, many elements and facilities have to be balanced across a server to deliver maximum throughput. As with previous generations of systems based on POWER processors, the design philosophy for POWER7 processor-based systems is one of system-wide balance in which the POWER7 processor plays an important role.

Each POWER7 processor has two integrated memory controllers in the chip. Industry standard DDR3 Registered DIMM (RDIMM) technology is used to increase reliability, speed, and density of memory subsystems. The functional minimum memory configuration for the machine is 4 GB (2 x 2 GB) but that is not sufficient for reasonable production use of the machine. It is recommended for the IBM Flex System p460 Compute Node a minimum of 32 GB of memory, with 32 x 16 GB DIMMs the maximum memory configurable is 512 GB.

The p460 has an onboard SAS controller that can manage up to two, non-hot-pluggable internal drives. Even though the p460 is a full-width server, it has the same storage options as the p260 and the p24L.

The type of local drives used impacts the form factor of your memory DIMMs. If HDDs are chosen, then only very-low-profile (VLP) DIMMs can be used because of internal spacing. There is not enough room for the 2.5-inch drives to be used with low-profile (LP) DIMMs (currently the 2 GB and 16 GB sizes). Verify your memory choice to make sure it is compatible with the local storage configuration. The use of SSDs does not have the same limitation, and LP DIMMs can be used with SSDs.

The p460 System Management is the same as the p260 and p24L POWER compute nodes.

I/O adapter slots

The networking subsystem of the IBM Flex System Enterprise Chassis has been designed to provide increased bandwidth and flexibility. The new design also allows for more ports on the available expansion adapters, which will allow for greater flexibility and efficiency with your system's design.

There are four I/O adapter slots on the IBM Flex System p460 Compute Node. The I/O adapters for the p460 have their own connector that plugs into the IBM Flex System Enterprise Chassis midplane. There is no onboard network capability in the Power Systems compute nodes other than the Flexible Service Processor (FSP) NIC interface.

The I/O is controlled by four P7-IOC I/O controller hub chips. This provides additional flexibility when assigning resources within Virtual I/O Server (VIOS) to specific Virtual Machine/LPARs.

1.5 I/O modules

The Enterprise Chassis can accommodate a total of four I/O modules which are installed in vertical orientation into the rear of the chassis, as shown in Figure 1-13, where the four modules at the back of the chassis with the bays numbered. In addition to the two types of switches listed in Table 1-1 on page 6, there are alternative I/O modules that provide external connectivity, as well as connecting internally to each of the nodes within the chassis. They can be either Switch or Pass through modules with a potential to support other types in the future. These models can be ordered in "build to order" IBM Flex Systems solutions.




Figure 1-13 IBM Flex System Enterprise Chassis with I/O module bays numbered

If a node has a two port integrated LAN on Motherboard (LOM) as standard, Module 1 and 2 are connected to it. If an I/O adapter is installed in the nodes I/O expansion bay 1, then Module 1 and 2 would be connected to this. Module 3 and 4 connect to the I/O adapter that is installed within I/O expansion bay 2 on the node. See Figure 1-14.

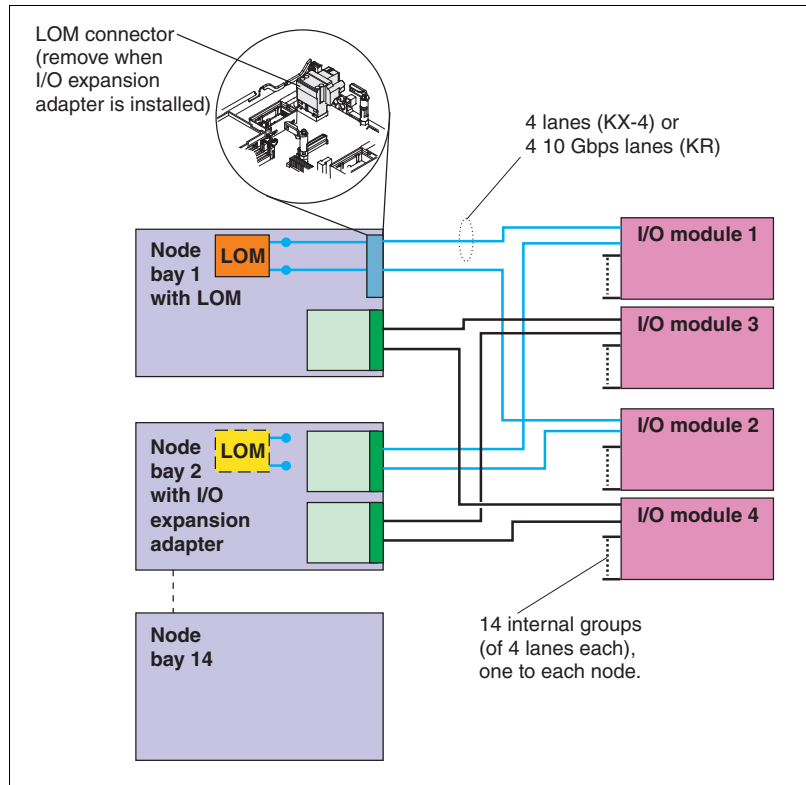


Figure 1-14 LOM, I/O adapter and switch module connection for node bays

The node in Bay 1 in Figure 1-14 shows that when shipped with a LOM, the LOM connector provides the link from the node motherboard to the midplane. Some nodes do not ship with LOM.

If required, this LOM connector can be removed and an I/O expansion adapter installed in its place. It is shown on the node in Bay 2 in Figure 1-14.

1.5.1 IBM Flex System Fabric CN4093 10 Gb Converged Scalable Switch

The IBM Flex System Fabric CN4093 10 Gb Converged Scalable Switch provides support for L2 and L3 switching, Converged Enhanced Ethernet (PFC, ETS, DCBX), Fibre Channel over Ethernet (FCoE), NPV Gateway, and Full Fabric Fibre Channel Forwarder (FCF).

The switch has the following major components:

- ▶ 42 10 Gb Ethernet internal ports and twenty-two external ports. External ports are arranged as two (small form-factor pluggable plus) SFP+ ports
- ▶ 12 SFP+ Omni Ports
- ▶ 2 Quad Small Form-Factor Pluggable Plus (QSFP+) ports.
- ▶ Each Omni Port is capable of running in 10 Gb Ethernet or 4/8 Gb FC mode with auto-negotiation capability.

- ▶ Support for Converged Enhanced Ethernet (CEE) and Fibre Channel over Ethernet (FCoE) over all Ethernet ports including Omni Ports (in Ethernet mode).
- ▶ Support for the Full Fabric FCF (Fibre Channel Forwarder) and NPV gateway.
- ▶ Support for full fabric FC services including Name Server, hardware-based Zoning.
- ▶ Support for IBM vNIC (virtual network interface card) Virtual Fabric Adapter with Single Root I/O Virtualization (SR-IOV) capability.

The 10 Gb Ethernet switch supports single compute node port capability (14 ports). Dual compute node port capability (28 ports) and triple compute node port capability (42 ports) are available with optional licenses.

The base model of this scalable switch provides the following features:

- ▶ 14 internal 10 Gb Ethernet/FCoE ports
- ▶ 2 external 1 Gb/10 Gb Ethernet/FCoE ports
- ▶ 6 external flexible ports, usable for either 10 Gb Ethernet/FCoE or 4/8 Gb Fibre Channel.

With the optional licensing for pay-as-you-grow scalability, you can easily and cost-effectively enable additional internal 10 Gb Ethernet/FCoE ports, external 10 Gb/40 Gb Ethernet/FCoE ports and external flexible ports, usable for either 10 Gb Ethernet/FCoE or 4/8 Gb Fibre Channel.

For switch management, access can be provided through the following connections:

- ▶ A SSHv2/Telnet connection to the embedded command-line interface (CLI)
- ▶ A terminal emulation program connection to the serial port interface
- ▶ A Web browser-based interface (https/http) connection to the switch

1.5.2 IBM Flex System Fabric EN4093 and EN4093R 10 Gb Scalable Switch

The IBM Flex System Fabric EN4093 and EN4093R 10 Gb Scalable Switches are 10 Gb 64-port upgradable midrange to high-end switch module, offering Layer 2/3 switching designed to install within the I/O module bays of the Enterprise Chassis. The switch has the following features:

- ▶ Up to 42 internal 10 Gb ports
- ▶ Up to 14 external 10 Gb uplink ports (SFP+ connectors)
- ▶ Up to 2 external 40 Gb uplink ports (QSFP+ connectors)

The switch is considered particularly suited for these needs:

- ▶ Building a 10 Gb infrastructure
- ▶ Implementing a virtualized environment
- ▶ Investment protection for 40 Gb uplinks
- ▶ TCO reduction, improving performance, while maintaining high levels of availability and security
- ▶ Oversubscription avoidance (traffic from multiple internal ports attempting to pass through a lower quantity of external ports, leading to congestion and performance impact)

The rear of the switch has 14 SFP+ module ports and 2QSFP+ module ports. The QSFP+ ports can be used to provide either two 40 Gb uplinks or eight 10 Gb ports, using one of the supported QSFP+ to 4x 10 Gb SFP+ cables. This cable splits a single 40 Gb QSFP+ port into 4 SFP+ 10 Gb ports.

For management of the switch, a mini USB port and also an Ethernet management port are provided.

IBM Flex System Fabric EN4093R: IBM Flex System Fabric EN4093R's stacking capabilities simplify management for clients by stacking up to eight switches that share one IP address and one management interface. Support for Switch Partition (SPAR) allows clients to virtualize the switch with partitions that isolate communications for multitenancy environments.

For more information about the IBM Flex System Fabric EN4093 and EN4093R 10 Gb Scalable Switches, see *IBM Flex System Fabric EN4093 and EN4093R 10Gb Scalable Switches*, TIPS0864.

1.5.3 IBM Flex System EN4091 10 Gb Ethernet Pass-thru

The EN4091 10Gb Ethernet Pass-thru module offers a 1 for 1 connection between a single node bay and an I/O module uplink. It has no management interface and can support both 1 Gb and 10 Gb dual-port adapters installed in the compute nodes. If quad-port adapters are installed in the compute nodes, only the first two ports will have access to the pass-thru module's ports.

The necessary 1 GbE or 10 GbE module (SFP, SFP+ or DAC) must also be installed in the external ports of the pass-thru, to support the speed wanted (1 Gb or 10 Gb) and medium (fiber optic or copper) for adapter ports on the compute nodes.

Four-port adapters: The EN4091 10 Gb Ethernet Pass-thru has only 14 internal ports. As a result, only two ports on each compute node are enabled, one for each of two pass-thru modules installed in the chassis. If four-port adapters are installed in the compute nodes, ports 3 and 4 on those adapters are not enabled.

For more information about the IBM Flex System EN4091 10 Gb Ethernet Pass-thru, see *IBM Flex System EN4091 10Gb Ethernet Pass-thru Module*, TIPS0865.

1.5.4 IBM Flex System EN2092 1 Gb Ethernet Scalable Switch

The EN2092 1 Gb Ethernet Switch provides support for L2/L3 switching and routing. The switch has:

- ▶ Up to 28 internal 1 Gb ports
- ▶ Up to 20 external 1 Gb ports (RJ45 connectors)
- ▶ Up to 4 external 10 Gb uplink ports (SFP+ connectors)

The switch comes standard with 14 internal and 10 external Gigabit Ethernet ports enabled. Further ports can be enabled, including the four external 10 Gb uplink ports.

For more information about the IBM Flex System EN2092 1 Gb Ethernet Scalable Switch, see *IBM Flex System EN2092 1Gb Ethernet Scalable Switch*, TIPS0861.

1.5.5 IBM Flex System FC5022 16 Gb SAN Scalable Switch

The IBM Flex System FC5022 16 Gb SAN Scalable Switch is a high-density, 48-port 16 Gbps Fibre Channel switch that is used in the Enterprise Chassis. The switch provides 28 internal ports to compute nodes by way of the midplane, and 20 external SFP+ ports. These SAN switch modules deliver an embedded option for IBM Flex System users deploying storage area networks in their enterprise. They offer end-to-end 16 Gb and 8 Gb connectivity.

The N_Port Virtualization mode streamlines the infrastructure by reducing the number of domains to manage while enabling the ability to add or move servers without impact to the SAN. Monitoring is simplified by an integrated management appliance, or clients using end-to-end Brocade SAN can leverage the Brocade management tools.

Two versions are available, a 12-port switch module and a 24-port switch with the Enterprise Switch Bundle (ESB) software. The port count can be applied to internal or external ports using a feature called Dynamic Ports on Demand (DPOD).

With DPOD, ports are licensed as they come online. With the FC5022 16Gb SAN Scalable Switch, the first 12 ports reporting (on a first-come, first-served basis) on boot-up are assigned licenses. These 12 ports can be any combination of external or internal Fibre Channel (FC) ports. After all licenses have been assigned, you can manually move those licenses from one port to another. As it is dynamic, no defined ports are reserved except ports 0 and 29. The FC5022 16Gb ESB Switch has the same behavior, the only difference is the number of ports.

For more information about the IBM Flex System FC5022 16 Gb SAN Scalable Switch, see *IBM Flex System FC5022 16Gb SAN Scalable Switches*, TIPS0870.

1.5.6 IBM Flex System FC3171 8 Gb SAN Switch

The IBM Flex System FC3171 8 Gb SAN Switch is a full-fabric Fibre Channel switch module that can be converted to a pass-thru module when configured in transparent mode. It can be done using the switch GUI or CLI and then the Module can be converted back to a full function SAN switch at some future date. The switch requires a reset when turning transparent mode on or off.

The I/O module has 14 internal ports and 6 external ports. All ports are licensed on the switch as there are no port licensing requirements.

On this switch, when in Full Fabric mode, access to all of the Fibre Channel Security features is provided. Security includes additional services available, such as Secure Socket Layer (SSL) and Secure Shell (SSH). In addition, RADIUS servers can be used for device and user authentication. After SSL/SSH is enabled, then the Security features are available to be configured. This allows the SAN administrator to configure which devices are allowed to login to the Full Fabric Switch module, by creating security sets with security groups. They are configured on a per switch basis. The security features are not available when in pass-thru mode.

The switch can be configured either by command line or by QuickTools:

- ▶ **Command Line:** Access the switch by the console port through the IBM Flex System Chassis Management Module or through the Ethernet Port. This method requires a basic understanding of the CLI commands.
- ▶ **QuickTools:** Requires a current version of the JRE on your workstation before pointing a web browser to the switch's IP address. The IP Address of the switch must be configured. QuickTools does not require a license and code is included.

For more information about the IBM Flex System FC3171 8 Gb SAN Switch, see *IBM Flex System FC3171 8Gb SAN Switch and Pass-thru*, TIPS0866.

1.5.7 IBM Flex System FC3171 8 Gb SAN Pass-thru

The IBM Flex System FC3171 8 Gb SAN Pass-thru I/O module is an 8 Gbps Fibre Channel pass-thru SAN module that has 14 internal ports and six external ports. It is shipped with all ports enabled.

Tip: If there is a potential future requirement to enable full fabric capability, then this switch should not be purchased and instead the FC3171 8Gb SAN Switch should be considered.

The FC3171 8 Gb SAN Pass-thru can be configured using either command line or QuickTools.

- ▶ **Command Line:** Access the module by the console port through the IBM Flex System Chassis Management Module or through the Ethernet Port. This method requires a basic understanding of the CLI commands.
- ▶ **QuickTools:** Requires a current version of the JRE on your workstation before pointing a web browser to the modules IP address. The IP Address of the module must be configured. QuickTools does not require a license and code is included.

For more information about the IBM Flex System FC3171 8 Gb SAN Pass-thru, see *IBM Flex System FC3171 8Gb SAN Switch and Pass-thru*, TIPS0866.

1.5.8 IBM Flex System IB6131 InfiniBand Switch

The IBM Flex System IB6131 InfiniBand Switch is a 32 port InfiniBand switch. It has 18 FDR/QDR (56/40 Gbps) external ports and 14 FDR/QDR (56/40 Gbps) internal ports for connections to nodes. This switch ships standard with QDR and can be upgraded to FDR.

Running the MLNX-OS, this switch has one external 1 Gb management port and a mini USB Serial port for updating software and debug use, along with InfiniBand internal and external ports.

The switch has fourteen internal QDR links and eighteen CX4 uplink ports. All ports are enabled. The switch can be upgraded to FDR speed (56 Gbps) by the Feature On Demand (FOD) process.

Tip: InfiniBand is not a supported protocol for IBM Flex System V7000 Storage Node or the Storwize V7000.

For more information about the IBM Flex System IB6131 InfiniBand Switch, see *IBM Flex System IB6131 InfiniBand Switch*, TIPS0871.

1.6 Introduction to IBM Flex System storage

Either the IBM Storwize V7000 or IBM Flex System V7000 Storage Node is an integrated part of the IBM PureFlex System, depending on the model. Figure 1-15 shows the IBM Flex System V7000 Storage Node where the left controller, called a cannister, is taken out. For the IBM Storwize V7000 products, the cannisters mount from the rear, whereas in the IBM Flex System V7000 Storage Node, the controllers mount from the front.



Figure 1-15 IBM Flex System V7000 Storage Node

The IBM Storwize V7000 product is described in detail in the *Implementing the IBM Storwize V7000 V6.3*, SG24-7938.

For more information about the IBM Flex System V7000 Storage Node, see Chapter 2, “Introduction to IBM Flex System V7000 Storage Node” on page 39.

1.6.1 IBM Storwize V7000 and IBM Flex System V7000 Storage Node

IBM Storwize V7000 and IBM Flex System V7000 Storage Node are virtualized storage systems designed to complement virtualized server environments. They provide unmatched performance, availability, advanced functions, and highly scalable capacity. IBM Storwize V7000 and IBM Flex System V7000 Storage Node are powerful disk systems that have been designed to be easy to use and enable rapid deployment without additional resources. IBM Storwize V7000 and IBM Flex System V7000 Storage Node support block workloads, whereas Storwize V7000 Unified (not covered in this book) consolidates block and file workloads into a single storage system for simplicity of management and reduced cost. IBM Storwize V7000 and IBM Flex System V7000 Storage Node support IBM Real-time Compression™ for extraordinary efficiency, enabling storage of up to five times as much active data in the same physical disk space¹.

IBM Storwize V7000 and IBM Flex System V7000 Storage Node offer greater efficiency and flexibility through built-in solid state drive (SSD) optimization and thin provisioning technologies. IBM Storwize V7000 and IBM Flex System V7000 Storage Node advanced functions also enable non-disruptive migration of data from existing storage, simplifying implementation and minimizing disruption to users. Finally, these systems also enable you to virtualize and reuse existing disk systems, supporting a greater potential return on investment (ROI).

¹ Compression data based on IBM measurements. Compression rates vary by data type and content

IBM Flex System V7000 Storage Node is designed to integrate into the IBM PureFlex System or IBM Flex System to enable extremely rapid storage deployment and breakthrough management simplicity. This new class of storage system combines no-compromise design along with virtualization, efficiency, and performance capabilities of IBM Storwize V7000. It helps simplify and speed PureFlex System and IBM Flex System infrastructure deployment with superior server and storage management integration to automate and streamline provisioning and to help organizations achieve greater responsiveness to business needs while reducing costs.

For more information about IBM Real-time Compression in relation to IBM SAN Volume Controller and IBM Storwize V7000, see *Real-time Compression in SAN Volume Controller and Storwize V7000*, REDP-4859.

Highlights

Here are the highlights of IBM Storwize V7000 and IBM Flex System V7000 Storage Node:

- ▶ Delivers sophisticated enterprise-class storage function for businesses of all sizes
- ▶ Supports your growing business requirements while controlling costs
- ▶ Provides up to 200 percent performance improvement with automatic migration to high-performing Solid State Drives
- ▶ IBM Storwize V7000 and IBM Flex System V7000 Storage Node enable storing up to five times as much active data in the same physical disk space using IBM Real-time Compression²
- ▶ Enables near-continuous availability of applications through dynamic migration
- ▶ Supports faster and more efficient data copies for online backup, testing or data mining
- ▶ Offers flexible server and storage management with easy to use GUI for block and file storage management

IBM Storwize V7000 and IBM Flex System V7000 Storage Node are powerful block storage systems that combine hardware and software components to provide a single point of control to help support improved storage efficiency. By enabling virtualization, consolidation, and tiering in business of all sizes, it is designed to improve application availability and resource utilization. The system offers easy-to-use, efficient, and cost-effective management capabilities for both new and existing storage resources in your IT infrastructure.

Enhancing access with Easy Tier

IBM Easy Tier® provides automatic migration of frequently accessed data to high performing Solid State Drives (SSDs), enhancing usage efficiencies. Operating at a fine granularity, the Easy Tier function automatically repositions pieces of the data to the appropriate class of drives based on I/O patterns and drive characteristics with no further administrative interaction.

Easy Tier makes it easy and economical to deploy SSDs in your environment. A hybrid pool of storage capacity is created containing two tiers: SSD and Hard Disk Drive (HDD). The busiest portions of volumes are identified and automatically relocated to high-performance SSDs. Remaining data can take advantage of higher capacity, price-optimized drives for the best customer value. Volumes in an SSD-managed or HDD-managed disk group are monitored and can be managed automatically or manually by moving hot extents to SSD and cold extents to HDD.

² IBM lab measurements

With an online database workload, Easy Tier improved throughput up to 200 percent and reduced transaction response time by up to 30 percent compared to a configuration using only HDD.²

Extraordinary storage efficiency

IBM Storwize V7000 and IBM Flex System V7000 Storage Node combine a variety of IBM technologies including thin provisioning, automated tiering, storage virtualization, Real-time Compression, clustering, replication, multi-protocol support, and a next-generation graphical user interface (GUI). Together, these technologies enable IBM Storwize V7000 and IBM Flex System V7000 Storage Node to deliver extraordinary levels of storage efficiency.

Newest of these technologies is IBM Real-time Compression, which is designed to improve efficiency by compressing data as much as 80 percent, enabling you to store up to five times as much data in the same physical disk space. Unlike other approaches to compression, IBM Real-time Compression is designed to be used with active primary data such as production databases and e-mail applications, which dramatically expands the range of candidate data that can benefit from compression. As its name implies, IBM Real-time Compression operates immediately as data is written to disk so no space is wasted storing uncompressed data awaiting post-processing.

The benefits of using IBM Real-time Compression together with other efficiency technologies are very significant and include reduced acquisition cost (because less hardware is required), reduced rack space, and lower power and cooling costs throughout the lifetime of the system. When combined with external storage virtualization, Real-time Compression can significantly enhance the usable capacity of your existing storage systems, extending their useful life even further.

IBM Real-time Compression is available for IBM Storwize V7000 and IBM Flex System V7000 Storage Node.

Avoiding disruptions with dynamic migration

IBM Storwize V7000 and IBM Flex System V7000 Storage Node use virtualization technology to help insulate host applications from physical storage changes. This ability can help enable applications to run without disruption while you make changes to your storage infrastructure. Your applications keep running so you can stay open for business.

Moving data is one of the most common causes of planned downtime. IBM Storwize V7000 and IBM Flex System V7000 Storage Node include a dynamic data migration function that is designed to move data from existing block storage into the new system or between arrays in a IBM Storwize V7000 and IBM Flex System V7000 Storage Node, while maintaining access to the data. The data migration function might be used, for example, when replacing older storage with newer storage, as part of load balancing work or when moving data in a tiered storage infrastructure.

Using the dynamic migration capabilities can provide efficiency and business value. Dynamic migration can speed time-to-value from weeks or months to days, minimize downtime for migration, eliminate the cost of add-on migration tools, and can help avoid penalties and additional maintenance charges for lease extensions. The result can be real cost savings to your business.

Foundation for cloud deployments

Improving efficiency and delivering a flexible, responsive IT infrastructure are essential requirements for any cloud deployment. Key technologies for delivering this infrastructure include virtualization, consolidation, and automation.

With their virtualized storage design and tight affinity with technologies such as IBM PowerVM and VMware, IBM Storwize V7000, IBM Flex System V7000 Storage Node, and IBM Storwize V7000 Unified are the ideal complement for virtualized servers that are at the heart of cloud deployments.

IBM Storwize V7000 and IBM Flex System V7000 Storage Node help enable consolidation of multiple storage systems for greater efficiency. With IBM Storwize V7000 and IBM Flex System V7000 Storage Node, clustered systems drive the value of consolidation much further, and IBM Real-time Compression improves the cost effectiveness even more. Automated tiering technologies such as Easy Tier, IBM Active Cloud Engine™, and Tivoli software help make the best use of the storage resources available.

Protecting data with replication services

IBM Storwize V7000 and IBM Flex System V7000 Storage Node support block data, while Storwize V7000 Unified supports both file and block data in the same system with replication functions optimized for the specific needs of each type of data.

Integrated management

IBM Storwize V7000 and IBM Flex System V7000 Storage Node provide a tiered approach to management designed to meet the diverse needs of different organizations. The systems' management interface is designed to give administrators intuitive control of these systems and provides a single integrated approach for managing both block and file storage requirements in the same system.

For organizations looking to manage both physical and virtual server infrastructures and the storage they consume (including provisioning and monitoring for higher availability, operational efficiency and infrastructure planning), IBM Storwize V7000, IBM Flex System V7000 Storage Node are integrated with IBM Systems Director Storage Control and IBM Flex System Manager™. A single administrator can manage and operate IBM servers (IBM System x, IBM Power Systems, IBM BladeCenter and IBM PureFlex System) along with networking infrastructure and IBM storage from a single management panel.

High-performance SSD support

For applications that demand high disk speed and quick access to data, IBM provides support for SSDs in 200 and 400 GB 2.5-inch E-MLC (enterprise-grade multilevel cell) capacity. For ultra-high-performance requirements, IBM Storwize V7000 and IBM Flex System V7000 Storage Node can be configured with only SSDs for up to 96 TB of physical capacity in a single system (384 TB in a clustered system), enabling scale-out high performance SSD support.

External storage virtualization

External storage virtualization is the ability of IBM Storwize V7000 and IBM Flex System V7000 Storage Node to manage capacity in other disk systems. When IBM Storwize V7000 and IBM Flex System V7000 Storage Node virtualize a disk system, its capacity becomes part of the IBM Storwize V7000 and IBM Flex System V7000 Storage Node systems and is managed in the same way as capacity on internal drives. Capacity in external disk systems inherits all the functional richness and ease-of-use of IBM Storwize V7000 and IBM Flex System V7000 Storage Node including advanced replication, thin provisioning, Real-time Compression, and Easy Tier. Virtualizing external storage helps improve administrator productivity and boost storage utilization while also enhancing and extending the value of an existing storage asset.

External Storage Virtualization: External Virtualization is only supported over FC or FCOE interface at this time. ISCSI is not supported for this function.

For more information about External storage virtualization see Chapter 7, “Storage Migration Wizard” on page 265.

1.6.2 Benefits and value proposition

With IBM Storwize V7000 and IBM Flex System V7000 Storage Node, you get the following benefits:

- ▶ Simplified management and intuitive Graphical User Interface (GUI) aid in rapid implementation and deployment.
- ▶ Virtualization of existing storage infrastructure improves administrator productivity.
- ▶ You can improve space utilization up to 33 - 50 percent and up to 75 percent less capacity needed with IBM FlashCopy® snapshots.
- ▶ With IBM Storwize V7000 and IBM Flex System V7000 Storage Node IBM Real-time Compression is designed to improve efficiency by storing up to five times as much active primary data in the same physical disk space. By significantly reducing storage requirements, you can keep up to five times more information online, use the improved efficiency to reduce storage costs, or achieve a combination of greater capacity and reduced cost.
- ▶ Storage performance is increased up to 200 percent using Easy Tier technology.
- ▶ Dynamic migration helps decreasing migration times from weeks or months to days, eliminate the cost of add-on migration tools and provides continuous availability of applications by eliminating downtime.
- ▶ Thin provisioning allows you to purchase only the disk capacity needed.
- ▶ With IBM Storwize V7000 and IBM Flex System V7000 Storage Node, clustered systems support the needs of growing business while enabling you to buy additional hardware only as needed.

1.6.3 Data Protection features

The following Data Protection features are supported with With IBM Storwize V7000 and IBM Flex System V7000 Storage Node:

- ▶ Volume Mirroring allows a volume to remain online even when the storage pool backing it becomes inaccessible. The mirror copy is local to the system.
- ▶ Metro Mirror is a type of Remote Copy that creates a synchronous copy of data from a master volume to an auxiliary volume. The mirror copy is placed on a remote system and is an exact copy of the primary volume.
- ▶ Global Mirror provides an asynchronous copy, which means that the secondary volume is not an exact match of the primary volume at every point in time. The Global Mirror function provides the same function as Metro Mirror Remote Copy without requiring the hosts to wait for the full round-trip delay of the long-distance link. Like Metro Mirror, the mirror copy is placed on a remote system.

Mirroring limitations: Global (long distance of greater than 300 Km) mirroring is supported only over an FC SAN infrastructure. Local Mirroring (less than 300 Km) is supported over either an FC or FCoE SAN infrastructure. Mirroring over iSCSI fabric is not supported at this time. Consult System Storage Interoperability Center for supported configurations.

- ▶ Remote replication functions create exact copies of your data at remote locations to help you stay up and running in case of an emergency.
- ▶ FlashCopy and snapshot functions create instant copies of data to minimize data loss.

1.7 External storage

In addition to the IBM Flex System V7000 Storage Node, the IBM Flex System Enterprise Chassis offers several possibilities for integration into storage infrastructure, such as Fibre Channel, iSCSI, and Converged Enhanced Ethernet.

There are several options for attaching external storage systems to Enterprise Chassis, including these possibilities:

- ▶ Storage area networks (SANs) based on Fibre Channel technologies
- ▶ SANs based on iSCSI
- ▶ Converged Networks based on 10 Gb Converged Enhanced Ethernet (CEE)

1.7.1 Storage products

Fibre Channel-based SANs are the most common and advanced design of external storage infrastructure. They provide high levels of performance, availability, redundancy, and scalability. However, the cost of implementing FC SANs will be higher in comparison with CEE or iSCSI. The major components of almost every FC SAN include server's HBAs, FC switches, FC storage servers, FC tape devices, and optical cables for connecting these devices to each other.

iSCSI-based SANs provide all the benefits of centralized shared storage in terms of storage consolidation and adequate levels of performance, but use traditional IP-based Ethernet networks instead of expensive optical cabling. iSCSI SANs consist of server hardware iSCSI adapters or software iSCSI initiators, traditional network components such as switches, routers, and so forth, and storage servers with an iSCSI interface, such as IBM System Storage DS3500 or IBM N Series.

Converged Networks are capable of carrying both SAN and LAN types of traffic over the same physical infrastructure. Such consolidation allows you to decrease costs and increase efficiency in building, maintaining, operating, and managing of the networking infrastructure.

iSCSI, FC-based SANs, and Converged Networks can be used for diskless solutions to provide greater levels of utilization, availability, and cost effectiveness.

At the time of writing, the following IBM System Storage products are supported with the Enterprise Chassis:

- ▶ IBM Storwize V7000
- ▶ IBM XIV® Storage System series
- ▶ IBM System Storage DS8000® series
- ▶ IBM System Storage DS5000 series

- ▶ IBM System Storage DS3000 series
- ▶ IBM System Storage N series
- ▶ IBM System Storage TS3500 Tape Library
- ▶ IBM System Storage TS3310 Tape Library
- ▶ IBM System Storage TS3100 Tape Library

For the latest support matrices for storage products, see the storage vendors' interoperability guides. IBM storage products can be referenced in the IBM System Storage Interoperability Center (SSIC):

<http://www.ibm.com/systems/support/storage/ssic/interoperability.wss>

For the purpose of this book, we limit the information to IBM Storwize V7000. For more information about the other supported IBM System Storage products, see *IBM PureFlex System and IBM Flex System Products and Technology*, SG24-7984.

1.7.2 IBM Storwize V7000

IBM Storwize V7000 is an innovative storage offering that delivers essential storage efficiency technologies and exceptional ease of use and performance, all integrated into a compact, modular design. IBM Flex System V7000 Storage Node architecture is the same as that of the IBM Storwize V7000 and is managed from the IBM Flex System Chassis Management Module or IBM Flex System Manager node. IBM Storwize V7000 is considered external storage from a IBM PureFlex System perspective.

There are four levels of integration of Storwize V7000 with IBM Flex System as shown in Table 1-3.

Table 1-3 Levels of integration

Level	Integration
Starting Level	IBM Flex System Single Point of Management
Higher Level	<ul style="list-style-type: none"> ▶ Datacenter Management ▶ IBM Flex System Manager Storage Control
Detailed Level	<ul style="list-style-type: none"> ▶ Data Management ▶ Storwize V7000 Storage User GUI
Upgrade Level	<ul style="list-style-type: none"> ▶ Datacenter Productivity

For further information about IBM Storwize V7000, see the Redbooks publication, *Implementing the IBM Storwize V7000 V6.3*, SG24-7938, as well as the following website:

http://www.ibm.com/systems/storage/disk/storwize_v7000/overview.html

1.8 Useful websites

At the time this book is being written, IBM Flex System V7000 Storage Node is pre-leased and support pages are not yet available. Function and features of the IBM Storwize V7000 and IBM Flex System V7000 Storage Node are almost identical. Some useful links for the IBM Storwize V7000 can be used as an alternative:

- ▶ IBM PureSystems entry pages:

<http://www.ibm.com/ibm/puresystems/us/en/index.html>

- ▶ IBM PureSystems Redbooks pages:
<http://www.redbooks.ibm.com/portals/puresystems>
- ▶ IBM Storwize V7000 and IBM Flex System Redbooks pages:
<http://www.redbooks.ibm.com/cgi-bin/searchsite.cgi?query=v7000>
- ▶ IBM Flex System Information Center:
<http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp>
- ▶ IBM Flex System V7000 Storage Node Information Center:
http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.4939.doc/site_product_page.html
- ▶ IBM Flex System V7000 Storage Node - Learning Modules:
http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.4939.learning/site_learning_tutorials.html
- ▶ IBM Flex System V7000 Storage Node Support home:
http://www.ibm.com/support/entry/portal/overview/hardware/puresystems/pureflex_system/storage_node/flex_system_v7000
- ▶ IBM Storwize V7000 Information Center:
<http://publib.boulder.ibm.com/infocenter/storwize/ic/index.jsp>
- ▶ The IBM Storwize V7000 Support page is at the following website:
<http://www.ibm.com/storage/support/storwize/v7000>
- ▶ The IBM Storwize V7000 home page is at the following website:
http://www-03.ibm.com/systems/storage/news/center/storwize_v7000/index.html
- ▶ The IBM Storwize V7000 Supported hardware list:
<http://www-01.ibm.com/support/docview.wss?uid=ssg1S1003908>
- ▶ The IBM Storwize V7000 Configuration Limit and Restrictions are at the following website:
<http://www-01.ibm.com/support/docview.wss?uid=ssg1S1003902>
- ▶ IBM Flex System V7000 Storage Node Installation Guide (GC27-4203):
http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.4939.doc/site_bkmap_quickinstbk.pdf
- ▶ IBM Storwize V7000 Expansion Enclosure Installation Guide, Machine type 2076 (GC27-4234):
http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.4939.doc/tbrd_bkmap_install_2076.pdf
- ▶ IBM Flex System V7000 Storage Node Troubleshooting, Recovery, and Maintenance Guide (GC27-4205):
http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.4939.doc/site_bkmap_trblshoot.pdf



Introduction to IBM Flex System V7000 Storage Node

This chapter introduces the IBM Flex System V7000 Storage Node and the enclosures and capabilities on which it is based. We describe in detail the controller and the expansion enclosures that make up the hardware of IBM Flex System V7000 Storage Node and point out the differences between them. We present the concepts of virtualized storage and show how it works with the IBM Flex System V7000 Storage Node, as well as briefly describing the many software features and capabilities that are available with this environment.

2.1 IBM Flex System V7000 Storage Node overview

When virtualizing external storage arrays, an IBM Flex System V7000 Storage Node can provide up to 32 PB of usable capacity. An IBM Flex System V7000 Storage Node supports a range of external disk systems similar to what the IBM Storwize V7000 system supports today. A control enclosure contains two control canisters; an expansion enclosure contains two expansion canisters. Both of these enclosures can contain up to 24 disk drives of the 2.5 inch form factor.

The IBM Flex System V7000 Storage Node is a modular storage system designed to fit into the IBM Flex System Enterprise chassis. When purchased in the IBM PureFlex configurations, the IBM Flex System V7000 Storage Node is configured from the factory into the hardware solution purchased. If, however, the configuration wanted is not offered in the predefined offerings, then a “Build to Order” configuration is designed to meet your needs.

The IBM Flex System V7000 Storage Node includes the capability to virtualize its own internal storage in the same manner as the IBM Storwize V7000 does. The IBM Flex System V7000 Storage Node is built upon the software base of the IBM Storwize V7000, which uses technology from the IBM System Storage SAN Volume Controller (SVC) for virtualization and the advanced functions of the IBM System Storage DS8000 family for its RAID configurations of the internal disks, and the highly flexible graphical user interface (GUI) of the IBM XIV Storage Subsystem for management.

The IBM Flex System V7000 Storage Node provides a number of configuration options that are aimed at simplifying the implementation process. It also includes automated instruction steps, called *Directed Maintenance Procedures (DMP)*, to assist in resolving any events that might occur. The IBM Flex System V7000 Storage Node is a clusterable, scalable, storage system, and an external virtualization device.

The IBM Flex System V7000 Storage Node is designed to be a scalable internal storage system to support the compute nodes of the IBM Flex System environment. It will contain a control enclosure that contains a pair of clustered node canisters and accommodates up to twenty-four 2.5-inch disk drives within the enclosure. Each control enclosure can additionally attach a maximum of two IBM Flex System V7000 Expansion Enclosures that can reside in the IBM Flex System Enterprise chassis with it.

Optionally, up to nine IBM Storwize V7000 Expansion Enclosures can be installed externally; however, a total of no more than nine expansion enclosures using either IBM Flex System V7000 internal Expansion Enclosures (maximum 2), IBM Storwize V7000 external Expansion Enclosures (maximum 9), or any combination, are supported. With the configuration of a single IBM Flex System V7000 Control Enclosure using twenty-four 1 TB 2.5" disks attached to nine IBM Storwize V7000 Expansion enclosures, each using twelve 3 TB 3.5" disks, the system can manage a raw capacity of 348 TB.

The IBM Flex System V7000 Storage Node can contain up to four control enclosures in a cluster configuration, each supporting the full configuration described, resulting in a maximum total raw capacity of 1392 TB.

The control enclosure can support the addition of up to nine IBM Storwize V7000 expansion enclosures connected externally. The additional expansions allow for a total of 240 disk drives; or for a maximum capacity of 348 TB supported per control enclosure.

Tip: For maximum capacity numbers, the 7026-212 expansion enclosure and 3 TB SAS Nearline drives must be used.

Figure 2-1 shows a representation of the IBM virtual storage environment.

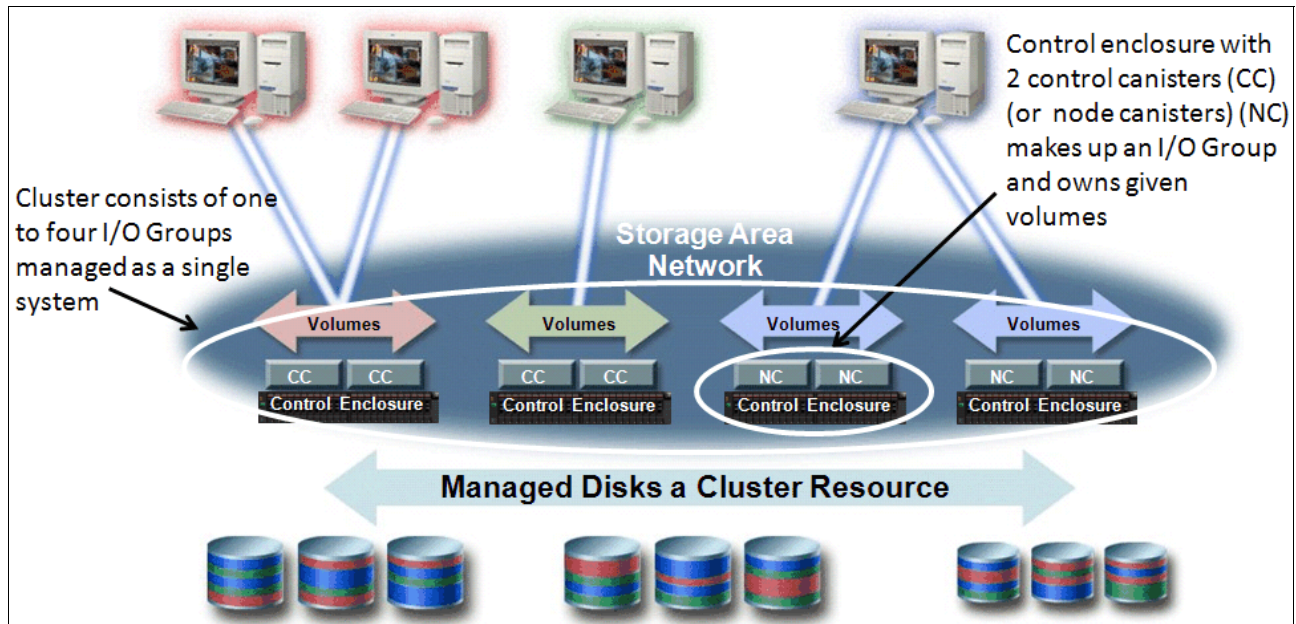


Figure 2-1 IBM virtual storage environment

2.2 IBM Flex System V7000 Storage Node terminology

The IBM Flex System V7000 Storage Node introduces some new terminology, which is defined in Table 2-1. We also include the terms first introduced with the IBM SVC and the IBM Storwize V7000, and are important to understand the rest of the implementation procedures described in this publication.

Table 2-1 IBM Flex System V7000 Storage Node terminology

IBM Flex System V7000 Storage Node term	Definition
Chain	The SAS2 connections by which expansion enclosures are attached which provides redundant access to the drives that are inside the enclosures. Each IBM Flex System V7000 Storage Node control canister has one chain connection.
Clone	A copy of a volume on a server at a particular point in time. The contents of the copy can be customized while the contents of the original volume are preserved.
Control canister	A hardware unit that includes all the management and control hardware, fabric and service interfaces, and the SAS2 expansion port.
Control enclosure	A hardware unit chassis that inserts into the IBM Flex System that includes control canisters including backup batteries, and 24 drive slots. It is the initial building block for the IBM Flex System V7000 Storage Node.
Event	An occurrence that is significant to a task or system. Events can include completion or failure of an operation, a user action, or the change in the state of a process.

IBM Flex System V7000 Storage Node term	Definition
Expansion canister	A hardware unit that includes the serial-attached SCSI (SAS2) interface hardware that enables the control enclosure to use the drives of the expansion enclosure as well as other expansions to be daisy-chained on behind it.
Expansion enclosure	A hardware unit chassis that inserts into the IBM Flex System that includes expansion canisters, and 24 drive slots used for connecting additional internal capacity to the IBM Flex System V7000 Storage Control Enclosure.
External V7000 expansion	A 2076-212 or 2076-224 IBM Storwize V7000 expansion that is connected to the IBM Flex System V7000 Storage Control Enclosure by the SAS2 chain to provide additional storage capacity which resides outside of the IBM Flex System.
External Virtualized storage	Managed disks (MDisks) that are presented as logical drives by external storage systems that are to be attached to the IBM Flex System V7000 Storage Node for additional virtual capacity.
Host mapping	The process of controlling which hosts have access to specific volumes within a clustered system.
Internal storage	The storage that resides in the control and expansion enclosures or the IBM Storwize V7000 expansions connected through the SAS2 chain that make up the IBM Flex System V7000 Storage Node.
Lane	The name given for a single 6Gbps SAS2 PHY (channel). There are four lanes (PHY) that make up each SAS2 chain.
Managed disk (MDisk)	A component of a storage pool that is managed by a clustered system. An MDisk is either a RAID array created using the internal storage, or a Small Computer System Interface (SCSI) logical unit (LU) for external storage being virtualized. An MDisk is not visible to a host system on the storage area network.
PHY	A term used to define a single 6Gbps SAS lane. There are four PHYs in each SAS cable.
Quorum disk	A disk that contains a reserved area that is used exclusively for cluster management. The quorum disk is accessed when it is necessary to determine which half of the cluster continues to read and write data. Quorum disks can either be on an MDisks or internal drives.
Snapshot	An image backup type that consists of a point-in-time view of a volume.
Storage pool	A collection of storage capacity on mdisks that can be used to provide the capacity requirements for a volume.
Strand	The serial-attached SCSI (SAS) connectivity of a set of drives within multiple enclosures. The enclosures can be either the IBM Flex System V7000 Storage control enclosures or expansion enclosures, or the IBM Storwize V7000 expansion enclosures.
Thin provisioning or Thin provisioned	The ability to define a storage unit (full system, storage pool, or volume) with a logical capacity size that is larger than the physical capacity assigned to that storage unit.
Volume	As used with the IBM Flex System V7000 Storage Node environment, it is the virtually defined device created for use by the host or the IBM Flex System V7000 Storage Node cluster to store I/O data.

2.3 IBM Flex System V7000 Storage Node

The IBM Flex System V7000 Storage Node is based on two enclosure types; the IBM Flex System V7000 Control Enclosure and the IBM Flex System V7000 Expansion Enclosure. Both of these enclosures reside in a newly designed common chassis that fits into the IBM Flex System Enterprise chassis. This new enclosure requires a space equal to two high and double wide slots to be available for each internal enclosure. Figure 2-2 shows an IBM Flex System V7000 Storage Node that has been built with a control enclosure and one expansion enclosure.




Figure 2-2 IBM Flex System V7000 Storage Node

Both the control and expansion enclosures connect to the Flex System Enterprise chassis through the midplane interconnect for their power and internal control connections. The control enclosure (A) also connects to the IBM Flex System I/O modules and switches for host I/O and replication features through this midplane. The control enclosure also houses a pair of redundant control canisters along with their cache batteries for backup.

The expansion enclosure (B) uses the Serial Attached SCSI (SAS2) 6 Gbps chain connection on the front of the control and expansion canisters (C) for connecting the chain together for drive I/O and expansion control operations. The expansion enclosure houses a pair of expansion canisters instead of the control canisters through which it connects and manages the SAS chain connections to its disk drives. It also has a second SAS2 port through which it provides a connection for continuing the chain to additional expansions behind it.

An IBM Flex System V7000 Storage Node is mainly intended to be a scalable, internal storage system, to support the internal compute nodes of the IBM Flex System. When needed, it can be expanded in its capacity by attaching external IBM Storwize V7000 expansion enclosures to its SAS2 chain. Both the 2076-212 and the 2076-224 model of the Storwize V7000 expansions are supported.

The control enclosure can support a combination of up to nine expansion enclosures using a combination of internal IBM Flex System V7000 Expansion Enclosures (maximum 2) and external IBM Storwize V7000 Expansion Enclosures (maximum 9) connected through the SAS2 connection on each of the control canisters. With the additional expansions, you are capable of up to 240 disks drives; or a maximum capacity of 348 TB per control enclosure.

2.3.1 IBM Flex System V7000 Storage Node capabilities

For a complete and updated list of IBM Flex System V7000 Storage Node configuration limits and restrictions, see the following website:

<http://www-01.ibm.com/support/docview.wss?uid=ssg1S1004114>

2.3.2 IBM Flex System V7000 Storage Node functions

The following functions are available with IBM Flex System V7000 Storage Node:

- ▶ Thin provisioning (no license required):
Traditional fully allocated volumes allocate real physical disk capacity for an entire volume even if that capacity is never used. Thin-provisioned volumes allocate real physical disk capacity only when data is written to the logical volume.
- ▶ Volume mirroring (no license required):
Provides a single volume image to the attached host systems while maintaining pointers to two copies of data in separate storage pools. Copies can be on separate disk storage systems that are being virtualized. If one copy is failing, IBM Flex System V7000 Storage Node provides continuous data access by redirecting I/O to the remaining copy. When the copy becomes available, automatic re-synchronization occurs.
- ▶ FlashCopy (included with the base IBM Flex System V7000 Storage Node license):
Provides a volume level point-in-time copy function for any storage being virtualized by IBM Flex System V7000 Storage Node. This function is designed to create copies for backup, parallel processing, testing, and development, and have the copies available almost immediately.

IBM Flex System V7000 Storage Node includes the following FlashCopy functions:

- Full / Incremental copy:
This function copies only the changes from either the source or target data since the last FlashCopy operation and is designed to enable completion of point-in-time online backups much more quickly than using traditional FlashCopy.
- Multitarget FlashCopy:
IBM Flex System V7000 Storage Node supports copying of up to 256 target volumes from a single source volume. Each copy is managed by a unique mapping and, in general, each mapping acts independently and is not affected by other mappings sharing the source volume.
- Cascaded FlashCopy:
This function is used to create copies of copies and supports full, incremental, or nocopy operations.
- Reverse FlashCopy:
This function allows data from an earlier point-in-time copy to be restored with minimal disruption to the host.

- FlashCopy nocopy with thin provisioning:

This function provides a combination of using thin-provisioned volumes and FlashCopy together to help reduce disk space requirements when making copies. There are two variations of this option:

- Space-efficient source and target with background copy:

Copies only the allocated space.

- Space-efficient target with no background copy:

Copies only the space used for changes between the source and target and is generally referred to as “snapshots”.

This function can be used with multi-target, cascaded, and incremental FlashCopy.

- Consistency groups:

Consistency groups address the issue where application data is on multiple volumes. By placing the FlashCopy relationships into a consistency group, commands can be issued against all of the volumes in the group. This action enables a consistent point-in-time copy of all of the data, even though it might be on a physically separate volume.

FlashCopy mappings can be members of a consistency group, or they can be operated in a stand-alone manner, that is, not as part of a consistency group. FlashCopy commands can be issued to a FlashCopy consistency group, which affects all FlashCopy mappings in the consistency group, or to a single FlashCopy mapping if it is not part of a defined FlashCopy consistency group.

- ▶ Remote Copy feature:

Remote Copy is an optional licensed feature that is based on the number of enclosures that are being used at the smallest configuration location. See 2.6.3, “Remote Copy licensing” on page 74 for licensing details. Remote Copy provides for the capability to perform either Metro mirror or Global Mirror operations:

- Metro Mirror:

Provides a synchronous remote mirroring function up to approximately 300 km between sites. As the host I/O only completes after the data is cached at both locations, performance requirements might limit the practical distance. Metro Mirror is designed to provide fully synchronized copies at both sites with zero data loss after the initial copy is completed.

Metro Mirror can operate between multiple IBM Flex System V7000 Storage Node systems and is only supported on either FC or FCoE host interfaces.

- Global Mirror:

Provides long distance asynchronous remote mirroring function up to approximately 8,000 km between sites. With Global Mirror, the host I/O completes locally and the changed data is sent to the remote site later. This function is designed to maintain a consistent recoverable copy of data at the remote site, which lags behind the local site.

Global Mirror can operate between multiple IBM Flex System V7000 Storage Node systems and is only supported on FC host interfaces.

- ▶ Data Migration (no charge for temporary usage):

IBM Flex System V7000 Storage Node provides a data migration function that can be used to import external storage systems into the IBM Flex System V7000 Storage Node system.

This function allows you to accomplish the following tasks:

- Move volumes non disruptively onto a newly installed storage system
- Move volumes to rebalance a changed workload
- Migrate data from other back-end storage to IBM Flex System V7000 Storage Node managed storage

▶ IBM System Storage Easy Tier (no charge):

Provides a mechanism to seamlessly migrate hot spots to the most appropriate tier within the IBM Flex System V7000 Storage Node solution. This migration could be to internal drives within IBM Flex System V7000 Storage Node or to external storage systems that are virtualized by IBM Flex System V7000 Storage Node.

▶ Real-time Compression (RTC):

Provides for data compression using the IBM Random-Access Compression Engine (RACE), which can be performed on a per volume basis in real time on active primary workloads. RTC can provide as much as a 50% compression rate for data that is not already compressed. It can help with reducing the amount of capacity needed for storage which can help with delaying further growth purchases. RTC supports all storage that is attached to the IBM Flex System V7000 Storage Node whether internal, external, or external virtualized storage.

A compression evaluation tool called the IBM Comprestimator can be used to determine the value of using compression on a specific workload for your environment.

▶ External Storage Virtualization (Licensed per enclosure of the external storage subsystem):

With this feature, an external storage subsystem can be attached through the Fibre Channel or by FCoE to the IBM Flex System V7000 Storage Node. These devices cannot be presented through an iSCSI connection. The devices presented are treated as mdisks and can be mapped to storage pools for volume creation and management. After the storage from the external system is integrated into the IBM Flex System V7000 Storage Node and added to a storage pool, it is available to be virtualized and used by any of the features and functions of the IBM Flex System V7000 Storage Node.

External virtualization: External Storage Virtualization is supported only on FC and FCoE but not iSCSI.

2.3.3 IBM Flex System V7000 Storage Node licensing

IBM Flex System V7000 Storage Node might require the following licenses:

- ▶ Enclosure
- ▶ External Virtualization
- ▶ Remote Copy (Advanced Copy Services: Metro Mirror / Global Mirror)
- ▶ Real-time Compression (RTC)

Table 2-2 gives a summary of all the licenses that might be required.

Table 2-2 Licenses that might be required

License type	Unit	License name	License required?
Enclosure	Base+expansion Physical Enclosure Number	IBM Flex System V7000 Base Software	Yes, software license per enclosure.
External Virtualization	Physical Enclosure Number Of External Storage	IBM Flex System V7000 External Virtualization Software	Optional add-on feature Yes, software license per external storage enclosure.
Remote Copy	Physical Enclosure Number	IBM Flex System V7000 Remote Mirroring Software	Optional add-on feature Yes, software license per enclosure.
Real-time Compression (RTC)	Physical Enclosure Number	IBM Flex System V7000 Real-time Compression Software	Optional add-on feature Yes, software license per enclosure.
FlashCopy	N/A	N/A	No
Volume Mirroring	N/A	N/A	No
Thin Provisioning	N/A	N/A	No
Volume Migration	N/A	N/A	No
Easy Tier	N/A	N/A	No

Migration: If the IBM Flex System V7000 Storage Node is used as a general migration tool, then the appropriate External Virtualization licenses must be ordered. The only exception is that if you want to migrate existing data from external storage to IBM Flex System V7000 Storage Node permanently managed storage, you can temporarily configure your External Storage license for use up to 45 days. For a more-than-45-day migration requirement from external storage to the IBM Flex System V7000 Storage Node, the appropriate External Virtualization license must be ordered.

2.4 IBM Flex System V7000 Storage Node hardware

The IBM Flex System V7000 Storage Node solution is a modular storage system that is built to interface and reside within the IBM Flex System Enterprise Chassis. When sold in the IBM PureFlex system it will come with a four port FC host interface card (HIC) preinstalled in it and configured to the FC switches in the Flex Enterprise Chassis. For the “Build to Order” (BTO) solutions you can select the configuration of the interface capability you require for your specific environment.

Host connectivity to compute nodes is provided through optional Flex System V7000 control enclosure network cards that connect to the Flex System Enterprise Chassis midplane and its switch modules. Available options are as follows:

- ▶ 10 Gb Converged Network Adapter (CNA) 2 Port Daughter Card for FCoE and iSCSI fabric connections
- ▶ 8 Gb Fibre Channel (FC) 4 Port Daughter Card for Fibre Channel fabric connections.

The Flex System V7000 Storage Node has two available slots in each control canister for populating with host network adapters. The adapters in each of the node canisters must be of the same type in each of the slots; so the adapters are installed in pairs (one adapter per node canister) and the following adapter configurations are supported for the entire Flex System V7000 Storage Node:

- ▶ Two or four 10 Gb CNA network adapters
- ▶ Two 4-port 8 Gb FC network adapters
- ▶ Two 10 Gb CNA network adapters and two 4-port 8 Gb FC network adapters

The 2-port 10 Gb Ethernet network adapters (up to two per canister) are used for iSCSI host attachment and/or FCoE attachment.

The configuration of the host attachments on one control canister must match the configuration of the second.

There is a 6 Gbps SAS port on the front of the canister for connecting optional expansions enclosures.

2.4.1 Control canister

The control canister is responsible for the management of all the virtualization, RAID functions, advanced features and functions of the IBM Flex System V7000 Storage Node and all command and I/O to its internal drive slots and the expansions that it is connected to. The control canister is a Customer Replaceable Unit (CRU). Figure 2-3 is a picture of the control canister with its cover removed.




Figure 2-3 Components and board layout of the control canister

In Figure 2-3, the processor (A) is a quad core Jasper. There are also two DIMMs which make up the cache memory and the battery backup unit (B) which beside the control canister is also a Customer Replaceable Unit (CRU).

As described before, the two host network adapter card locations (C and D) can also be seen. These provide the connections to the Flex System Enterprise Chassis through the midplane and the switch modules. It is important to remember that both control canisters must be populated with the same type of host network adapter cards in each of these locations.

Adapter locations: The first network adapter (slot1) location can only be populated by a 2-port 10 Gbps Ethernet CNA network adapter. The second location can be populated by either a 2-port 10 Gbps Ethernet CNA network adapter or a 4 port 8 Gbps Fibre Channel network adapter.

Figure 2-4 is a logical block diagram of the components and flow of the control canister.

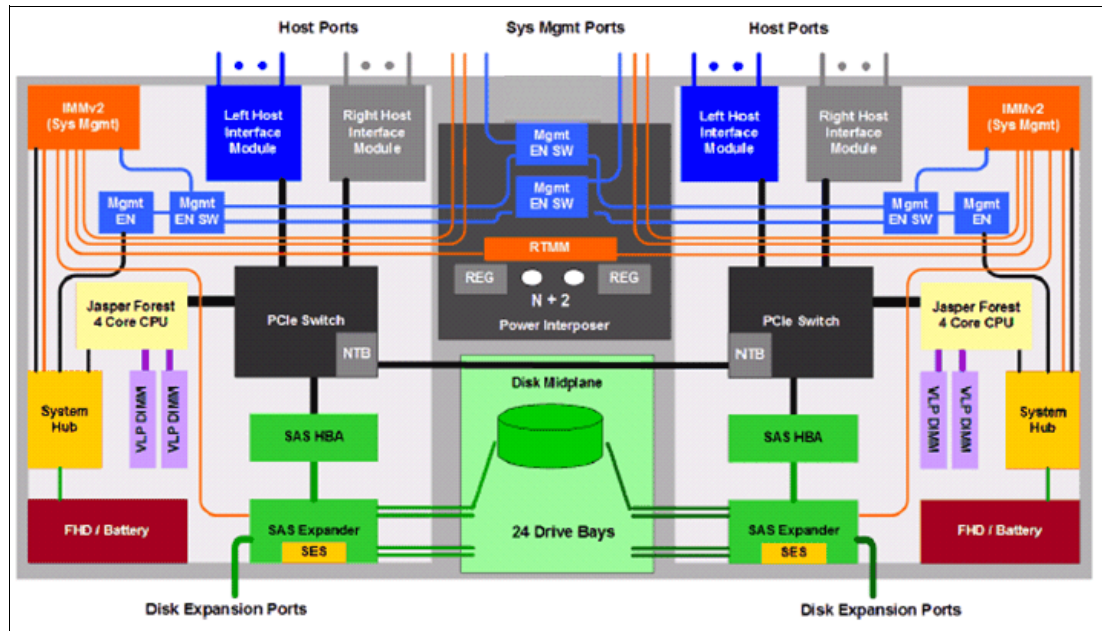


Figure 2-4 Logical block diagram for control canister

As shown in Figure 2-5, the control canister has only one SAS2 ports per canister for the connection of expansion enclosures to add capacity. There are also two USB connections available for support use to perform maintenance actions as required.

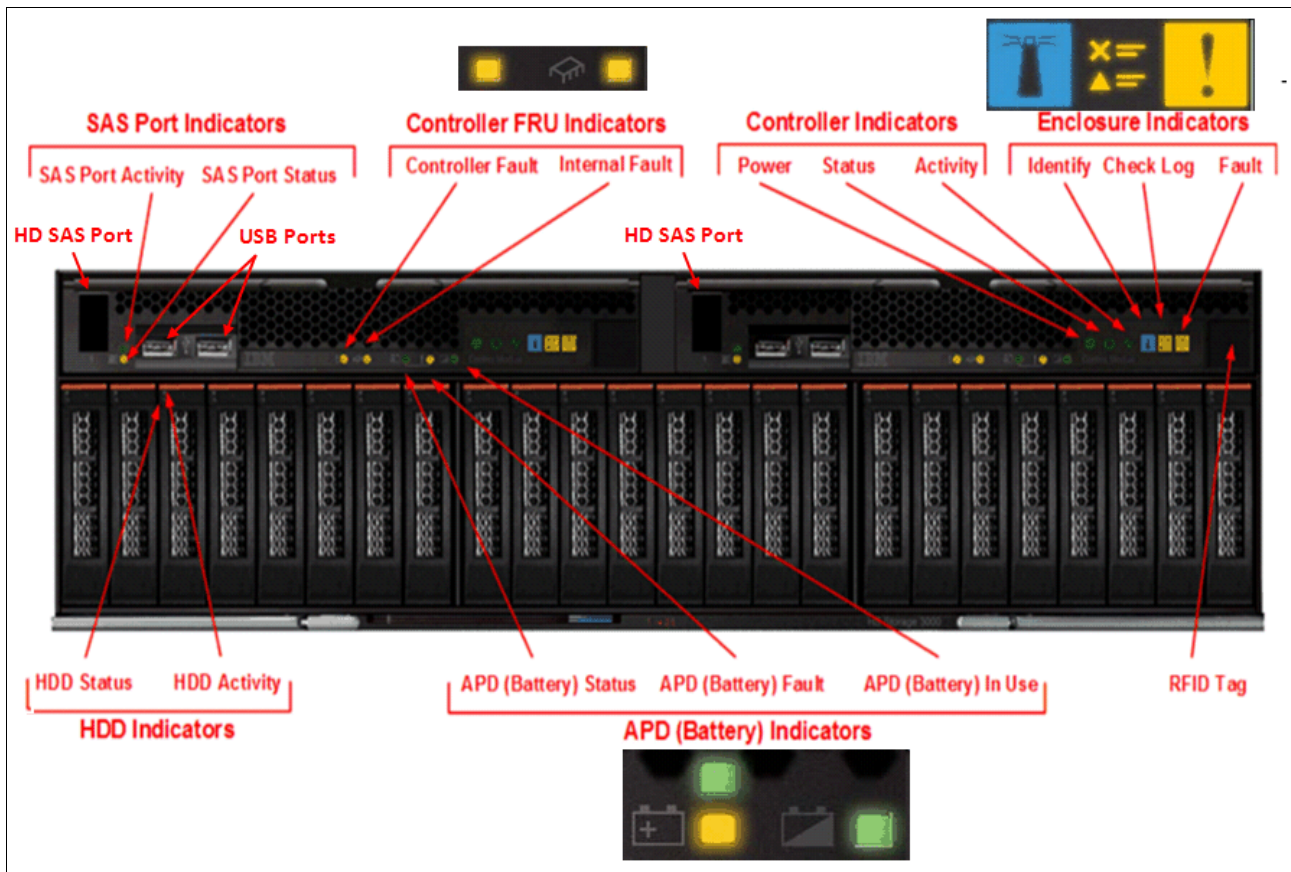


Figure 2-5 Control enclosure with connection ports and indicators

In Table 2-3, the LED indicators are defined for the control enclosure; some of them are also used on the expansion enclosure as well.

Table 2-3 Control enclosure LED description

LED Group	LED Name	Meaning
Enclosure Indicator The Enclosure Indicator LEDs on the righthand control canister are used as the primary enclosure LEDs. The left hand control canister LEDs are used when the righthand is not available.	Check Log	It is a software fault telling you to check the errorlog for details.
	Identify	It is used for identifying the selected enclosure.
	Fault	It indicates that a hardware fault has occurred.
Controller Indicators	Power	<ul style="list-style-type: none"> ▶ Slow Blink - Power available but Processor shutdown. ▶ Fast blink - doing POST. ▶ On solid - Powered up.
	Status	<ul style="list-style-type: none"> ▶ Off - not operational ▶ Solid - in a cluster ▶ Slow blinking - in cluster or service state ▶ Fast blink - is upgrading
	Activity	The Activity LED blinks to show there is I/O activity
Controller FRU Control canister hardware fault	Control	Hardware problem with the control canister.
	Internal FRU	Problem with an internal FRU like the network adapters.
Battery Indicators	In Use	Fast blink- System is shutting down on battery power.
	Status	<ul style="list-style-type: none"> ▶ Slow blink - Charging, but has enough power to boot the canister. ▶ Fast blink - Charging, and does not have enough power to boot the canister (Error 656) ▶ On solid - Fully charged.
	Fault	On - Hardware fault detected

2.4.2 Expansion canister

The expansion canister connects the expansion disks to the control canister using the SAS2 6 Gbps chain interface. This module also enables the daisy-chaining of additional expansions to be connected behind it to further expand the capacity of the system's storage. The expansion canister is also a Customer Replaceable Unit (CRU). Figure 2-6 shows a picture of the expansion canister with its cover removed.




Figure 2-6 Components and board layout of the expansion canister

As shown in Figure 2-6, the expansion canister does not contain a battery backup unit like the control canister does. It does have an additional SAS2 connection (A) on it to allow for the continuation of the chain to additional expansions. Figure 2-7 shows the SAS connection ports and the status indicators on the front of the expansion enclosures.




Figure 2-7 Expansion enclosure with connection ports and indicators

As the indicators on the expansion enclosure are a subset of the ones that are available on the control enclosure, Table 2-3 on page 51 provides the details of their definitions.

Figure 2-8 is a logical block diagram of the expansion canister.

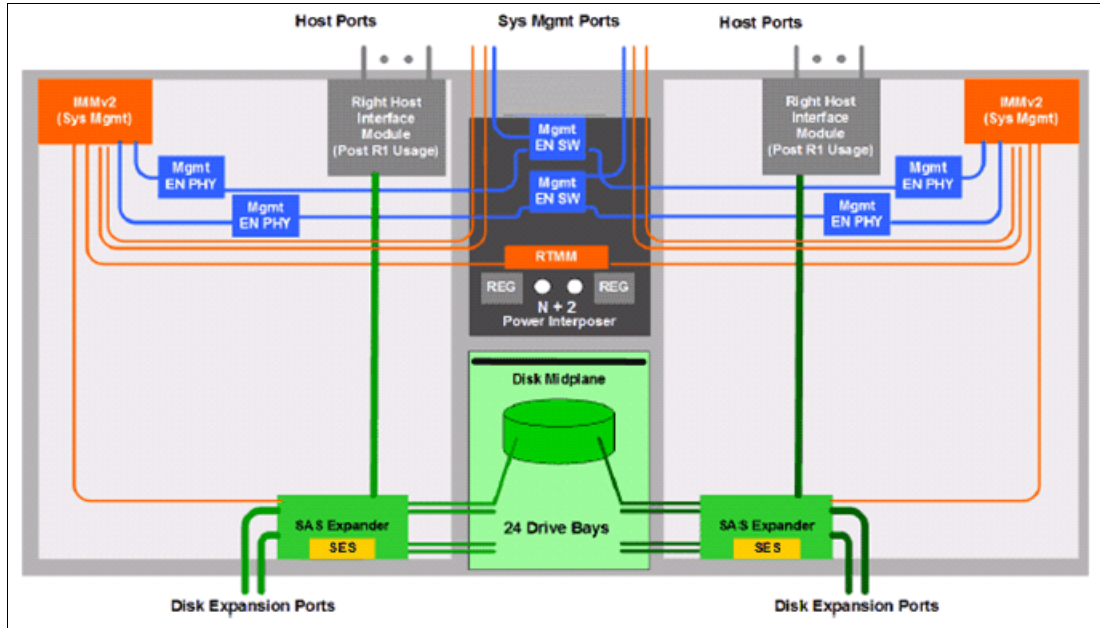


Figure 2-8 Logical block diagram for control canister

2.4.3 Supported disk drives

Both the IBM Flex System V7000 Control Enclosure and the IBM Flex System V7000 Expansion Enclosure support up to 24 2.5 inch disk drives in their enclosures. Table 2-4 shows all the possible drive types that can be used in the internal enclosure drive slots at the time of writing.

Table 2-4 IBM Flex System V7000 Storage Node internal supported drives

Drive capacity	Drive speed	Drive physical size
146 GB	15 KRPM SAS	2.5 inch
300 GB	15 KRPM SAS	2.5 inch
300 GB	10 KRPM SAS	2.5 inch
600 GB	10 KRPM SAS	2.5 inch
900 GB	10 KRPM SAS	2.5 inch
500 GB	7500 RPM NL SAS	2.5 inch
1 TB	7500 RPM NL SAS	2.5 inch
200 GB	SSD SAS	2.5 inch
400 GB	SSD SAS	2.5 inch

The disk drives connect to the Flex System Enterprise chassis through the midplane interconnect for their power. Also, in the control enclosure, the midplane interconnect is used for the internal control and I/O paths. The expansion enclosures use Serial Attached SCSI (SAS) connections on the front of the control and expansion canisters for disk I/O and control commands.

2.4.4 IBM Storwize V7000 expansion enclosure

The IBM Storwize V7000 expansion enclosure can be optionally attached to the IBM Flex System V7000 Storage Node for added capacity beyond that of the internal enclosures. These expansion enclosures contain two IBM Storwize V7000 *expansion canisters*, disk drives, and two power supplies. There are two models of the expansion enclosures with the 2076-212 providing 12 disk slots of the 3.5 inch form factor, and the 2076-224 providing 24 disk slots of the 2.5 inch drives form factor.

Figure 2-9 shows the components of the expansion enclosure.




Figure 2-9 Component side view Expansion enclosure

The expansion enclosure power supplies have a single power lead connector on the power supply unit. The PSU has an IEC C14 socket and the mains connection cable has a C13 plug. As shown in Figure 2-9, the PSU has one green status LED indicator (A) to show it is powered on and working properly.

Each expansion canister provides two SAS interfaces that are used to connect to either the IBM Flex System V7000 Storage Node control or expansion enclosures or to a preceding IBM Storwize V7000 expansion enclosure, as well as to connect any additional IBM Storwize V7000 expansion enclosure behind it. The ports are numbered 1 on the left and 2 on the right. SAS port 1 is the IN port and SAS port 2 is the OUT port. There is also a symbol printed at the ports to identify whether it is an IN or an OUT bound port.

Use of the SAS connector 1 is mandatory when installed, as the expansion enclosure must be attached to either a control enclosure or another expansion enclosure. SAS connector 2 is optional, as it is used to attach to the next additional expansion enclosure in the chain.

Each port connects four SAS physical links (PHYs) to the drives. As shown in Figure 2-9, there is an LED associated with each PHY in each port (eight LEDs in total per canister). The LEDs are green and grouped together next to the ports (B); for each port they are numbered 1 through 4. These LEDs indicate when there is activity on the PHY.

Figure 2-10 shows the front view of the 2076-212 enclosure.



Figure 2-10 IBM Storwize V7000 front view for 2076-212 enclosure

The drives are positioned in four columns of three horizontally mounted drive slots in the expansion enclosure. The drive slots are numbered 1 - 12, starting at the upper left and going from left to right, top to bottom.

Figure 2-11 shows the front view of the 2076-224 enclosure.



Figure 2-11 IBM Storwize V7000 front view for 2076-224 enclosure

The drives are positioned in one row of 24 vertically mounted row in the drive assemblies. The drive slots are numbered 1 - 24, starting from the left. (There is a vertical center drive bay molding between slots 12 and 13.)

Though the IBM Storwize V7000 2076-224 enclosure is a 2.5 inch, 24 drive slot chassis; the drives of this subsystem are not interchangeable with the IBM Flex System V7000 Storage Node drives. Both drives use a different carrier and contain a different product identifier in their code.

Therefore, the IBM Storwize V7000 expansions support their own drive types in their enclosures and should be created in their own configurations and storage pools when used with the IBM Flex System V7000 Storage Node.

IBM Flex System V7000 Storage Node enclosures currently support SSD, SAS, and Nearline SAS drive types. Each SAS drive has two ports (two PHYs) and I/O can be issued down both paths simultaneously.

For a list of supported drives for the IBM Storwize V7000 and other details, see the Information Center at the following website:

<http://pic.dhe.ibm.com/infocenter/storwize/ic/index.jsp>

2.4.5 SAS cabling requirements

The IBM Flex System V7000 Storage Node uses new smaller SAS2 cable connectors. These connectors are based on the high density (HD) mini SAS connectors.

For the connections between the IBM Flex System V7000 Control enclosure and the IBM Flex System V7000 Expansion enclosure in the chassis, you have to use these short cables (shown in Figure 2-2 on page 43) that can be ordered with your storage.

With the addition of external IBM Storwize V7000 expansion, there is an adapter that changes the HD mini SAS to the mini SAS connection that is on the expansion. Figure 2-12 shows the cabling scheme and the differences in the cable connections of the two expansions.

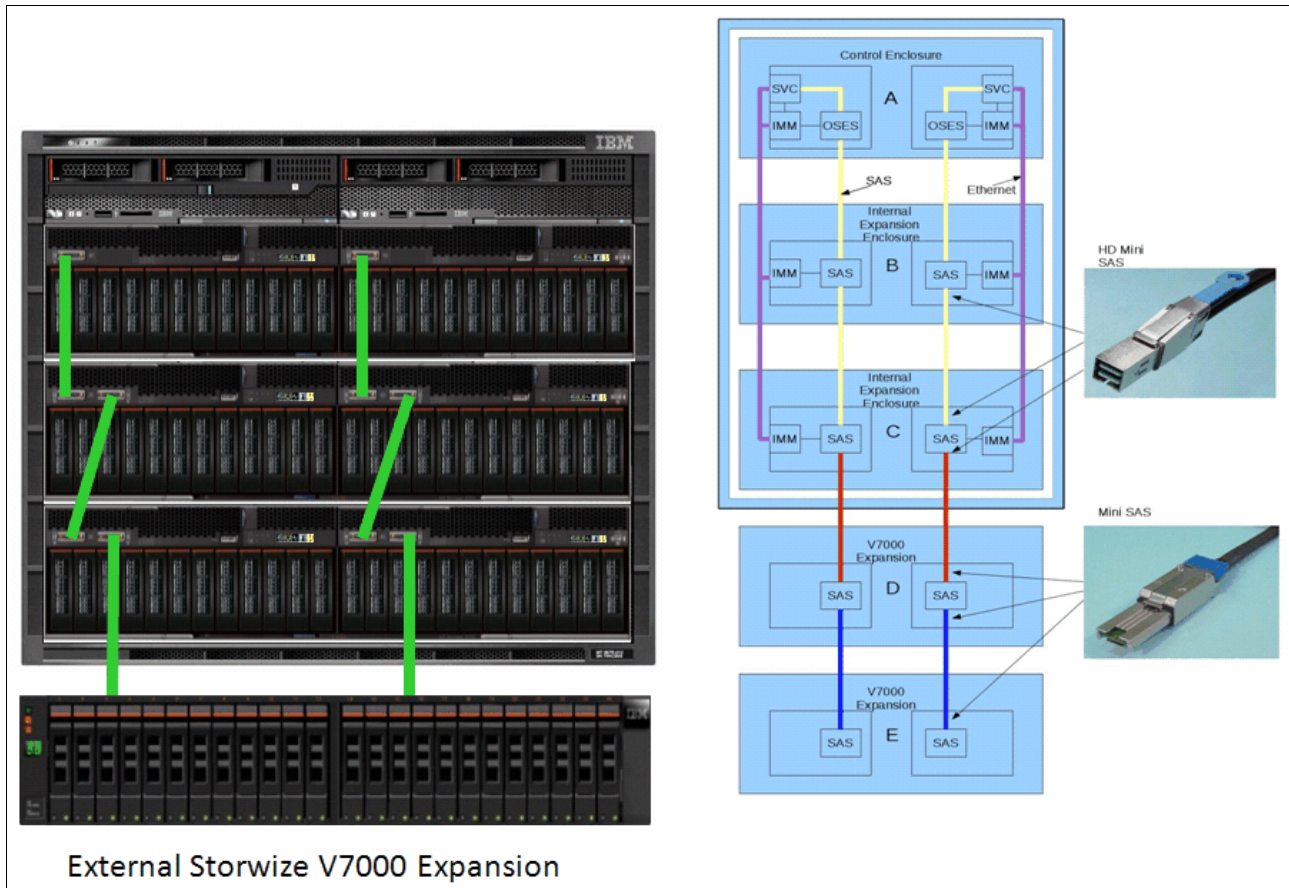


Figure 2-12 IBM Flex System V7000 Storage Node cabling internal SAS and external

Notice that the IBM Flex System V7000 Storage Node cabling is done directly from top to bottom down through its expansions, including any additional external expansions.

2.5 IBM Flex System V7000 Storage Node components

The IBM Flex System V7000 Storage Node is an integrated entry / midrange virtualization RAID storage to support and interface with the IBM Flex System. It brings with it the following benefits:

- ▶ Offers a single point of management for server, network, internal storage, and external storage, virtualized by the integrated IBM Flex System V7000 Storage Node.
- ▶ Gives improved access to critical data by keeping data close to the servers by minimizing switch and cable hops.
- ▶ Provides virtual servers fast shared storage access to better enable dynamic workload assignments.
- ▶ Moves the control of the virtual server storage over to the server administrators rather than in the hands of the SAN administrators.
- ▶ Uses the IBM Flex System Enterprise Chassis for lower solution costs, thus eliminating the need of external switches, cables, SFPs, fans, power supplies, for production environments.

Tip: Even though no external switches are required, the appropriate internal switch I/O modules must be available in order for the necessary host interfacing to work.

The IBM Flex System V7000 Storage Node consists of a set of drive enclosures. Control enclosures contain disk drives and two control canisters and form an I/O group for its management of additional internal and external storage. Expansion enclosures contain drives and are attached to control enclosures.

The simplest use of the IBM Flex System V7000 Storage Node is as a traditional RAID subsystem. The internal drives are configured into RAID arrays and virtual disks created from those arrays.

The IBM Flex System V7000 Storage Node can also be used to virtualize other storage controllers. An example of it is described in Chapter 7, “Storage Migration Wizard” on page 265.

The IBM Flex System V7000 Storage Node supports regular and solid-state drives. It can use IBM System Storage Easy Tier to automatically place volume hot spots on better-performing storage.

2.5.1 Hosts

A host system is an IBM Flex System compute node that is connected to IBM Flex System V7000 Storage Node through a Fibre Channel connection or through an Ethernet connection using either iSCSI or Fibre Channel over Ethernet (FCoE) connection through the switch modules of the IBM Flex System. At the time of writing, attachment of external hosts to the IBM Flex System V7000 Storage Node is not supported.

Hosts are defined to IBM Flex System V7000 Storage Node by identifying their worldwide port names (WWPNs) for Fibre Channel hosts. For iSCSI hosts, they are identified by using their iSCSI names. The iSCSI names can either be iSCSI qualified names (IQNs) or extended unique identifiers (EUIs).

2.5.2 Control canisters

IBM Flex System V7000 Storage Node can have two to eight hardware components called control canisters that provide the virtualization of internal and external volumes, and cache and copy services (Remote Copy) functions. Within an IBM Flex System V7000 Storage Node a pair of the control canisters are housed within a control enclosure. A clustered system consists of a one to four control enclosures.

One of the control canisters within the system is known as the configuration node and it is the canister that manages configuration activity for the clustered system. If this canister fails, the system nominates another canister to become the configuration node. During initial setup, the system will automatically select a canister for this role from the first control enclosure pair.

2.5.3 I/O groups

Within the IBM Flex System V7000 Storage Node, there can be one to four control enclosures that also are defined as *I/O groups*. The IBM Flex System V7000 Storage Node supports four control enclosures in the clustered system, which provides four I/O groups.

When a host server performs I/O to one of its volumes, all the I/Os for a specific volume are directed to the owning I/O group. Also, under normal conditions, the I/Os for that specific volume are always processed by the same canister within the I/O group.

Both canisters of the I/O group act as preferred nodes for their own specific subset of the total number of volumes that the I/O group presents to the host servers (a maximum of 2048 volumes). However, both canisters also act as the failover canister for its partner within the I/O group, so a canister can take over the I/O workload from its partner, if required.

In an IBM Flex System V7000 Storage Node environment, using active-active architecture, the I/O handling for a volume can be managed by both canisters of the I/O group. Therefore, it is mandatory for servers that are connected through Fibre Channel connectors to use multipath device drivers to be able to handle this capability.

The I/O groups are connected to the SAN so that all application servers accessing volumes from the I/O group have access to them. Up to 512 host server objects can be defined in two I/O groups.

Important: The active / active architecture provides ability to process I/Os for both control canisters and allows the application to continue running smoothly, even if the server has only one access route or path to the storage. This type of architecture eliminates the path / LUN thrashing that can exist with an active / passive architecture.

2.5.4 Clustered system

A clustered system consists of one to four control enclosures. All configuration, monitoring, and service tasks are performed at the system level and the configuration settings are replicated across all control canisters in the clustered system. To facilitate these tasks, one or two management IP addresses are set for the system.

There is a process provided to back up the system configuration data on to disk so that the clustered system can be restored in the event of a disaster. This method does not back up application data, only the IBM Flex System V7000 Storage Node configuration information.

System configuration backup: After backing up the system configuration, save the backup data on your hard disk (or at the least outside of the SAN). When you are unable to access the IBM Flex System V7000 Storage Node, you do not have access to the backup data if it is on the SAN.

For the purposes of remote data mirroring, two or more clustered systems (IBM Flex System V7000 Storage Nodes, Storwize V7000, or SAN Volume Controller systems starting from Version 6.4) must form a partnership before creating relationships between mirrored volumes.

Important: With IBM Flex System V7000 Storage Node 6.4, the **layer** parameter can only be changed by running **chsystem** using the CLI. The default is the **storage** layer, and you must change it to **replication** if you need to set up a copy services relationship between the IBM Flex System V7000 Storage Node with either the IBM Storwize V7000 or SAN Volume Controller.

As mentioned earlier, one canister is designated as the configuration node and it is the canister that activates the system IP address. If the configuration node fails, the system chooses a new canister as the configuration node and the new canister takes over the system IP addresses.

The system can be configured using either the IBM Flex System V7000 Storage Node management software the command-line interface (CLI),

2.5.5 RAID

The IBM Flex System V7000 Storage Node setup contains a number of internal disk drive objects known as candidate drives, but these drives cannot be directly added to storage pools. The drives need to be included in a Redundant Array of Independent Disks (RAID) grouping used for performance and to provide protection against the failure of individual drives.

These drives are referred to as members of the array. Each array has a RAID level. Different RAID levels provide different degrees of redundancy and performance, and have different restrictions regarding the number of members in the array.

IBM Flex System V7000 Storage Node supports hot spare drives. When an array member drive fails, the system automatically replaces the failed member with a hot spare drive and rebuilds the array to restore its redundancy (the exception being RAID 0). Candidate and spare drives can be manually exchanged with array members.

Each array has a set of goals that describe the wanted location and performance of each array. A sequence of drive failures and hot spare takeovers can leave an array unbalanced, that is, with members that do not match these goals. The system automatically rebalances such arrays when the appropriate drives are available.

The following RAID levels are available:

- ▶ RAID 0 (striping, no redundancy)
- ▶ RAID 1 (mirroring between two drives)
- ▶ RAID 5 (striping, can survive one drive fault)
- ▶ RAID 6 (striping, can survive two drive faults)
- ▶ RAID 10 (RAID 0 on top of RAID 1)

RAID 0 arrays stripe data across the drives. The system supports RAID 0 arrays with just one member, which is similar to traditional JBOD attach. RAID 0 arrays have no redundancy, so they do not support hot spare takeover or immediate exchange. A RAID 0 array can be formed by one to eight drives.

RAID 1 arrays stripe data over mirrored pairs of drives. A RAID 1 array mirrored pair is rebuilt independently. A RAID 1 array can be formed by two drives only.

RAID 5 arrays stripe data over the member drives with one parity strip on every stripe. RAID 5 arrays have single redundancy. The parity algorithm means that an array can tolerate no more than one member drive failure. A RAID 5 array can be formed by 3 to 16 drives.

RAID 6 arrays stripe data over the member drives with two parity stripes (known as the P-parity and the Q-parity) on every stripe. The two parity strips are calculated using different algorithms, which give the array double redundancy. A RAID 6 array can be formed by 5 to 16 drives.

RAID 10 arrays have single redundancy. Although they can tolerate one failure from every mirrored pair, they cannot tolerate two-disk failures. One member out of every pair can be rebuilding or missing at the same time. A RAID 10 array can be formed by 2 to 16 drives.

2.5.6 Managed disks

A managed disk (MDisk) refers to the unit of storage that IBM Flex System V7000 Storage Node virtualizes. This unit could be a logical volume from an external storage array presented to the IBM Flex System V7000 Storage Node or a RAID array created on internal drives, or an external Storwize V7000 expansion that is managed by the IBM Flex System V7000 Storage Node. The IBM Flex System V7000 Storage Node can allocate these MDisks into various storage pools for different usage or configuration needs.

An MDisk should not be visible to a host system on the storage area network, as it should only be zoned to the IBM Flex System V7000 Storage Node system.

An MDisk has four possible modes:

- ▶ **Array:**
Array mode MDisks are constructed from drives using the RAID function. Array MDisks are always associated with storage pools.
- ▶ **Unmanaged:**
Unmanaged MDisks are not being used by the system. This situation might occur when an MDisk is first imported into the system, for example.
- ▶ **Managed:**
Managed MDisks are assigned to a storage pool and provide extents so that volumes can use it.
- ▶ **Image:**
Image MDisks are assigned directly to a volume with a one-to-one mapping of extents between the MDisk and the volume. This situation is normally used when importing logical volumes into the clustered system that already have data on them, which ensures that the data is preserved as it is imported into the clustered system.

2.5.7 Quorum disks

A quorum disk is a managed disk (MDisk) that contains a reserved area for use exclusively by the system. In the IBM Flex System V7000 Storage Node, any managed disks can be considered as a quorum candidate. The clustered system uses quorum disks to break a tie when exactly half the control canisters in the system remain active after a SAN failure. The IBM Flex System V7000 Storage Node dynamically assigns which quorum drive will be the active member (DQ).

The diagram in Figure 2-13 shows a preferred quorum drive layout for a dual control enclosure clustered IBM Flex System V7000 Storage Node system.

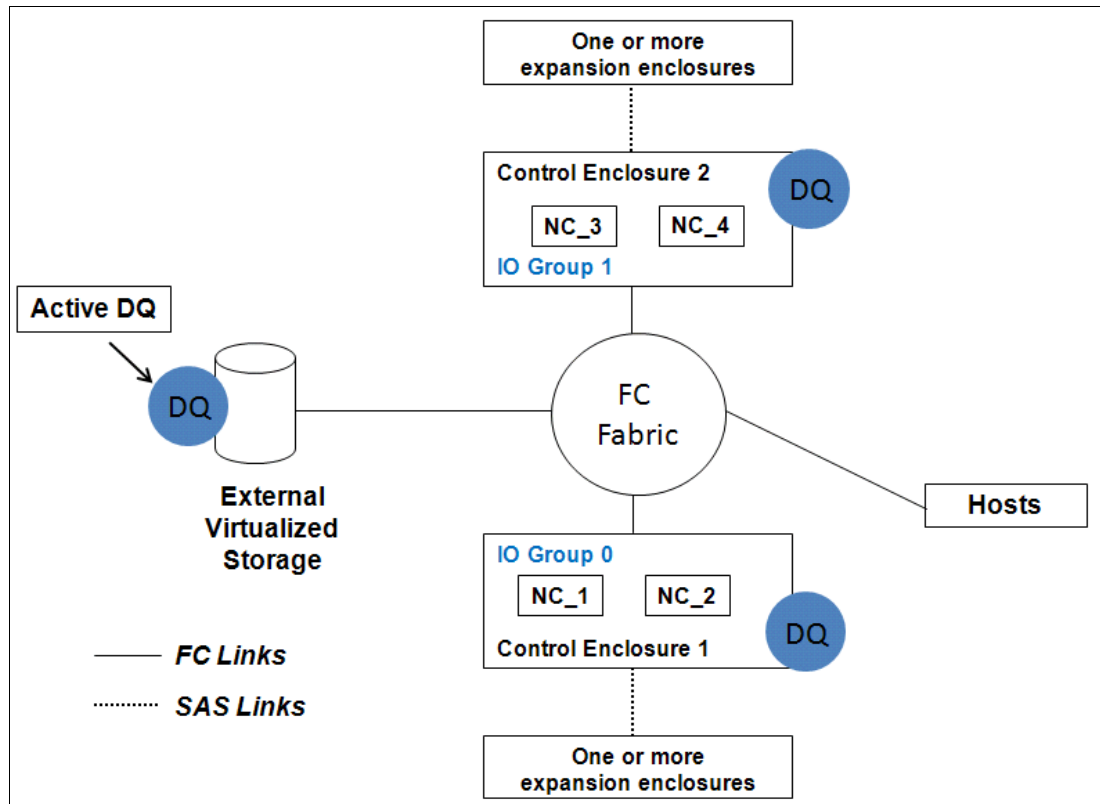


Figure 2-13 Preferred quorum drive layout

The clustered system automatically forms the quorum disk by taking a small amount of space from a managed disk (MDisk). It allocates space from up to three different MDisks for redundancy, although only one quorum disk is active.

If the environment has multiple storage systems, then to avoid the possibility of losing all of the quorum disks because of a failure of a single storage system, you should allocate the quorum disk on different storage systems. The preferred internal drives are the ones in the control enclosure and hot spares. However when an external virtualized storage system is being managed, it is preferred to have the active quorum disk located on it for better access when needed. It is possible to manage the quorum disks by using the CLI.

2.5.8 Storage pools

A storage pool is a collection of MDisks (up to 128) that are grouped together to provide capacity for the creation of the virtual volumes. All MDisks in the pool are split into extents with the same size. Volumes are then allocated out of the storage pool and are mapped to a host system.

Names: All object names must begin with an alphabetic character and cannot be numeric. The name can be a maximum of 63 characters. Valid characters are uppercase (A-Z), lowercase letters (a-z), digits (0 - 9), underscore (_), period (.), hyphen (-), and space; however, the names must not begin or end with a space.

MDisks can be added to a storage pool at any time to increase the capacity of the storage pool. MDisks can belong in only one storage pool and only MDisks in unmanaged mode can be added to the storage pool. When an MDisk is added to the storage pool, the mode changes from unmanaged to managed and vice versa when you remove it.

Each MDisk in the storage pool is divided into a number of extents. The size of the extent is selected by the administrator at creation time of the storage pool and cannot be changed later. The size of the extent ranges from 16 MB up to 8 GB.

The extent size has a direct impact on the maximum volume size and storage capacity of the clustered system. A system can manage 4 million (4 x 1024 x 1024) extents. For example, a system with a 16 MB extent size can manage up to 16 MB x 4 MB = 64 TB of storage.

The effect of extent size on the maximum volume size is shown in Table 2-5, and lists the extent size and the corresponding maximum clustered system size.

Table 2-5 Maximum volume capacity by extent size

Extent size	Maximum volume capacity for normal volumes (GB)
16	2048 (2 TB)
32	4096 (4 TB)
64	8192 (8 TB)
128	16384 (16 TB)
256	32768 (32 TB)
512	65536 (64 TB)
1024	131072 (128 TB)
2048	262144 (256 TB)
4096	528288 (512 TB)
8192	1056576 (1,024 TB)

The effect of extent size on the maximum clustered system capacity is shown in Table 2-6.

Table 2-6 Extent size and effect on clustered system capacity

Extent size	Maximum storage capacity of cluster
16 MB	64 TB
32 MB	128 TB
64 MB	256 TB
128 MB	512 TB
256 MB	1 PB
512 MB	2 PB
1024 MB	4 PB
2048 MB	8 PB
4096 MB	16 PB
8192 MB	32 PB

Use the same extent size for all storage pools in a storage system, when they are supporting volume migration between two storage pools. If the storage pool extent sizes are not the same, use volume mirroring to copy volumes between storage pools, as described in Chapter 10, “Volume mirroring and migration” on page 423.

For most IBM Flex System V7000 Storage Nodes, a maximum of capacity of 1 PB is sufficient; and an extent size value of 256 MB should be used.

Default extent size: The GUI of IBM Flex System V7000 Storage Node sets a default extent size value of 256 MB when you define a new storage pool.

A storage pool can have a threshold warning set that automatically issues a warning alert when the used capacity of the storage pool exceeds the set limit.

Single-tiered storage pool

MDisks that are used in a single-tiered storage pool need to have the following characteristics to prevent performance and other problems:

- ▶ They must have the same hardware characteristics, disk type, disk size and disk speeds (revolutions per minute (rpms)).
- ▶ They must have the same RAID type, RAID array size, and disk spindle count in the RAID grouping.
- ▶ The disk subsystems providing the MDisks must have similar characteristics, for example, maximum input/output operations per second (IOPS), response time, cache, and throughput.
- ▶ When possible, make all MDisks of the same size, and ensure that the MDisks provide the same number of extents per volume. If this configuration is not feasible, you need to check the distribution of the volumes' extents in that storage pool.

Multitiered storage pool

A multitiered storage pool has a mix of MDisks with more than one type of disk tier attribute, for example, a storage pool containing a mix of generic_hdd AND generic_ssd MDisks.

A multitiered storage pool contains MDisks with different characteristics as opposed to the single-tiered storage pool. However, each tier should have MDisks of the same size and MDisks that provide the same number of extents.

A multitiered storage pool is used to enable automatic migration of extents between disk tiers using the IBM Flex System V7000 Storage Node Easy Tier function. Figure 2-14 shows these components.

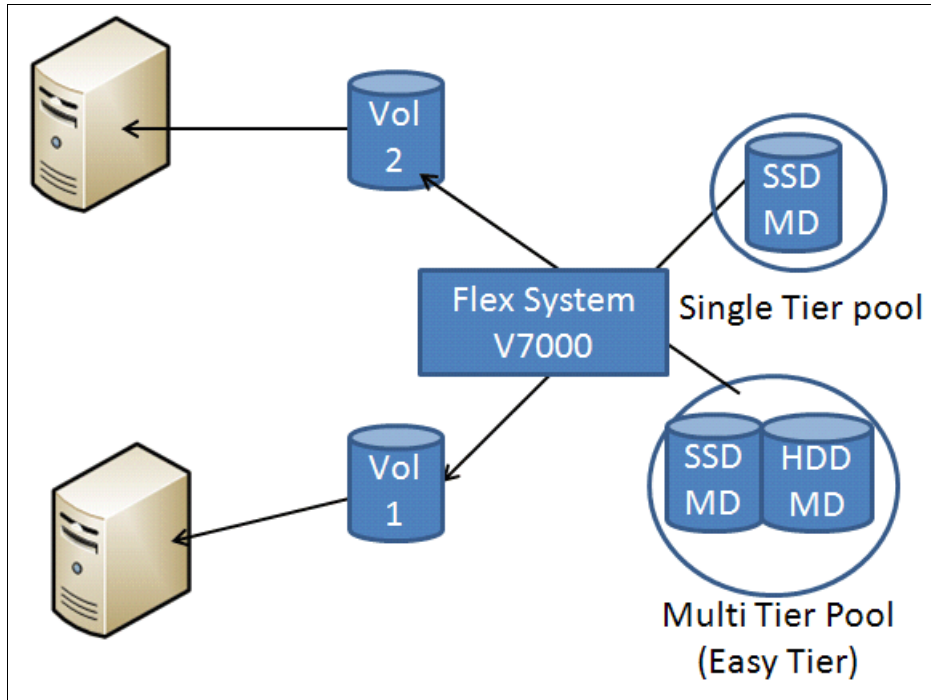


Figure 2-14 IBM Flex System V7000 Storage Node with single and multi-tier pools

2.5.9 Volumes

A volume is a virtual logical disk that is presented to a host system by the clustered system. In our virtualized environment, the host system has a volume mapped to it by IBM Flex System V7000 Storage Node. The IBM Flex System V7000 Storage Node translates this volume into a number of extents, which are allocated across MDisks. The advantage with storage virtualization is that the host is “decoupled” from the underlying storage, so the virtualization appliance can move the extents around without impacting the host system.

The host system cannot directly access the underlying MDisks in the same manner as it could access RAID arrays in a traditional storage environment.

There are three types of volumes:

► **Striped volume:**

A striped volume is allocated using one extent from each MDisk at a time in the storage pool. This process continues until the space required for the volume has been satisfied. It is also possible to supply a list of MDisks to use.

► **Sequential volume:**

A sequential volume is where the extents are allocated one after the other from one Mdisk. If there is not enough space on a single Mdisk, the creation of the sequential volume fails. If a volume is required to be expanded, the sequential volume is converted to a striped volume by policy when the expansion occurs.

Figure 2-15 shows examples of the striped and sequential volume types.

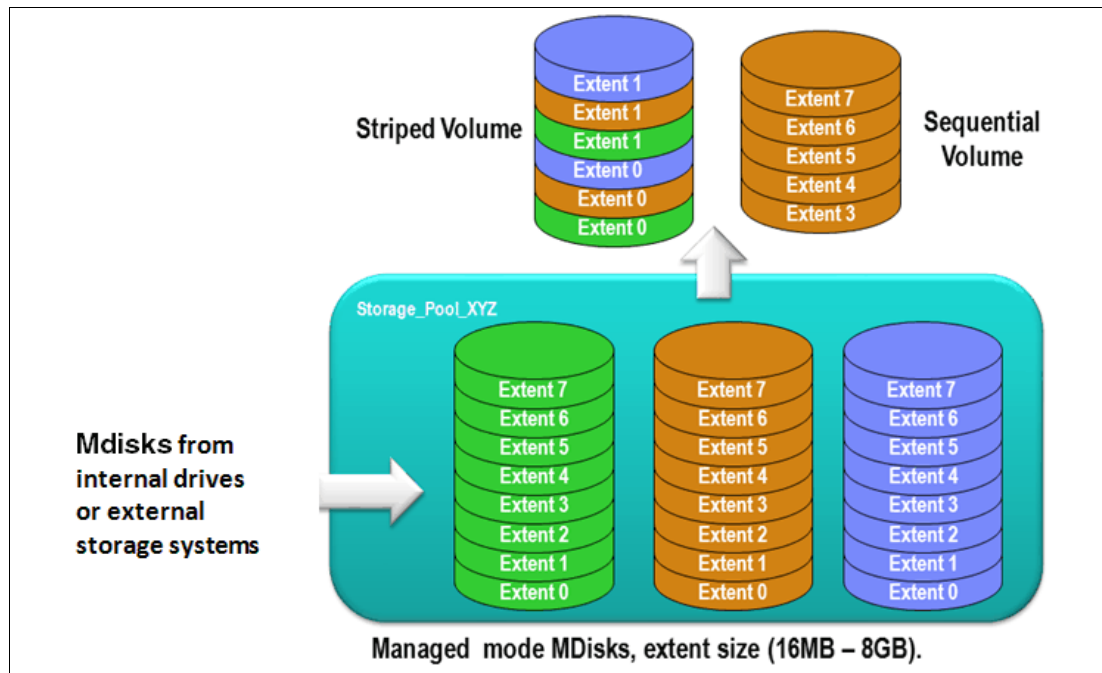


Figure 2-15 Striped and sequential volumes

► **Image mode:**

Image mode volumes are special volumes that have a direct relationship with one MDisk. They are used to migrate existing data into and out of the IBM Flex System V7000 Storage Node.

When the image mode volume is created, a direct mapping is made between extents that are on the MDisk and the extents that are on the volume. The logical block address (LBA) x on the MDisk is the same as the LBA x on the volume, which ensures that the data on the MDisk is preserved as it is brought into the clustered system (Figure 2-16).




Figure 2-16 Image mode volume

Some virtualization functions are not available for image mode volumes, so it is often useful to migrate the volume into a new storage pool. After it is migrated, the MDisk becomes a managed MDisk.

If you add an MDisk containing data to a storage pool, any data on the MDisk is lost. Ensure that you create image mode volumes from MDisks that contain data before adding MDisks to the storage pools.

2.5.10 Thin-provisioned volumes

Volumes can be configured to either be *thin provisioned* or *fully allocated*. A thin-provisioned volume behaves with respect to application reads and writes as though they were fully allocated. When a volume is created, the user specifies two capacities: the real capacity of the volume and its virtual capacity.

The real capacity determines the quantity of MDisk extents that are allocated for the volume. The virtual capacity is the capacity of the volume reported to IBM Flex System V7000 Storage Node and to the host servers.

The real capacity is used to store both the user data and the metadata for the thin-provisioned volume. The real capacity can be specified as an absolute value or a percentage of the virtual capacity.

The thin provisioning feature can be used on its own to create overallocated volumes, or it can be used with FlashCopy. Thin-provisioned volumes can be used with the mirrored volume feature as well.

A thin-provisioned volume can be configured to *autoexpand*, which causes IBM Flex System V7000 Storage Node to automatically expand the real capacity of a thin-provisioned volume as its real capacity is used. Autoexpand attempts to maintain a fixed amount of unused real capacity on the volume. This amount is known as the *contingency capacity*.

The contingency capacity is initially set to the real capacity that is assigned when the volume is created. If the user modifies the real capacity, the contingency capacity is reset to be the difference between the used capacity and real capacity.

A volume that is created with a zero contingency capacity goes offline as soon as it needs to expand. A volume with a non-zero contingency capacity stays online until it has been used up.

Autoexpand does not cause the real capacity to grow much beyond the virtual capacity. The real capacity can be manually expanded to more than the maximum that is required by the current virtual capacity, and the contingency capacity is recalculated.

To support the autoexpansion of thin-provisioned volumes, the storage pools from which they are allocated have a configurable warning capacity. When the used free capacity of the group exceeds the warning capacity, a warning is logged. For example, if a warning of 80% has been specified, the warning is logged when 20% of the free capacity remains.

A thin-provisioned volume can be converted to a fully allocated volume using *volume mirroring* (and vice versa).

2.5.11 Mirrored volumes

IBM Flex System V7000 Storage Node provides a function called *volume mirroring*, which enables a volume to have two physical copies. Each volume copy can belong to a different storage pool and can be on different managed physical storage systems, which can help provide a better level of high-availability to a solution.

When a host system issues a write to a mirrored volume, IBM Flex System V7000 Storage Node writes the data to both copies. When a host system issues a read to a mirrored volume, IBM Flex System V7000 Storage Node requests it from the primary copy. If one of the mirrored volume copies is temporarily unavailable, the IBM Flex System V7000 Storage Node automatically uses the alternate copy without any outage for the host system. When the mirrored volume copy is repaired, IBM Flex System V7000 Storage Node resynchronizes the data.

A mirrored volume can be converted into a non-mirrored volume by deleting one copy or by splitting one copy to create a new non-mirrored volume.

The mirrored volume copy can be any type; image, striped, sequential, and thin provisioned or not. The two copies can be different volume types.

Using mirrored volumes can also assist with migrating volumes between storage pools that have different extent sizes and can provide a mechanism to migrate fully allocated volumes to thin-provisioned volumes without any host outages.

You can change the time-out value setting to either *latency*, which prioritizes low host latency (default) or *redundancy*, which prioritizes redundancy (longer timeout).

Unmirrored volumes: An unmirrored volume can be migrated from one location to another by adding a second copy to the wanted destination, waiting for the two copies to synchronize, and then removing the original copy. This operation can be stopped at any time.

2.5.12 Easy Tier

Easy Tier is a performance function that automatically migrates or moves extents off a volume to, or from, SSD storage to HDD storage. Easy Tier monitors the host I/O activity and latency on the extent of all volumes with the Easy Tier function turned on in a multitiered storage pool over a 24-hour period. It then creates an extent migration plan based on this activity and then dynamically moves high activity or hot extents to a higher disk tier within the storage pool. It also moves extents' activity that has dropped off or cooled from the high tiered MDisk back to lower tiered MDisk.

It should be understood that Easy Tier does not do any operations unless the results will have a level of positive value that makes the activity worth performing. In an environment with low total workload on a volume, even if the volume has a specific hot spot, it might be judged to be too cool to justify the required operations to perform the move. Likewise, the downgrading of an extent will not take place until the space is needed for a hotter extent. Figure 2-17 shows the basic structure and concepts of this function.

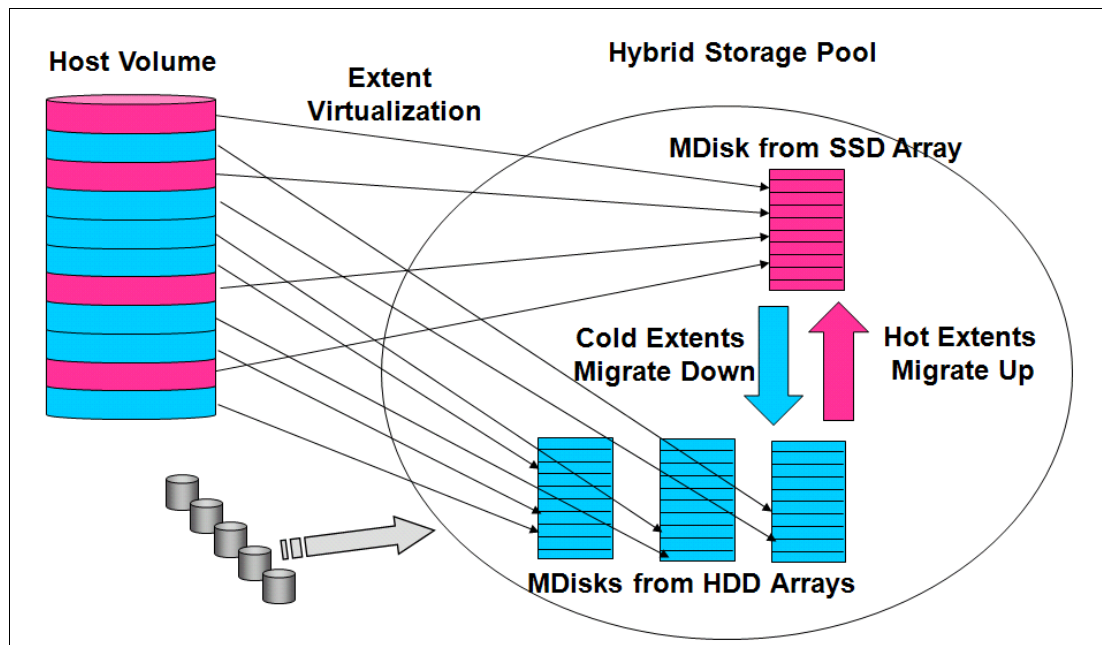


Figure 2-17 Easy Tier overview

The Easy Tier function can be turned on or off at the storage pool and volume level.

It is possible to demonstrate the potential benefit of Easy Tier in your environment before installing a solid-state disk. By turning on the Easy Tier function for a single level storage pool and the Easy Tier Function for the volumes within that pool, Easy Tier creates a migration report every 24 hours on the number of extents it would move if the pool was a multitiered pool. Easy Tier statistics measurement is enabled.

Using Easy Tier can make it more appropriate to use smaller storage pool extent sizes.

The usage statistics file can be offloaded from IBM Flex System V7000 Storage Node and then an IBM Storage Advisor Tool can be used to create a summary report from the data.

Contact your IBM representative or IBM Business Partner for more information about the Storage Advisor Tool. For more information about Easy Tier, see *Implementing the IBM Storwize V7000 V6.3*, SG24-7938.

2.5.13 Real-time Compression (RTC)

With the IBM Flex System V7000 Storage Node, there is a capability to create a volume as a compressed volume type. With this type of volume, storage capacity needs can be lowered by as much as half. IBM Flex System V7000 Storage Node RTC function is based off of the same proven Random-Access Compression Engine (RACE) as the IBM RTC Appliances.

With compression, storage growth can be curbed and need for additional storage purchases can be delayed and spread out over greater periods of time. RTC dynamically works with active workloads now compresses the data while it is being processed the first time.

To implement RTC, the volume must be created with compressed type selected. You cannot convert a volume from uncompressed to compressed after creation. However, a compressed volume can be a target of a volume mirror allowing the copying of the uncompressed volume to a compressed copy.

The RTC is a feature that is licensed on an enclosure basis for the IBM Flex System V7000 Storage Node.

RTC resource needs should be considered when planning for the use of compression. Resource requirements must come from someplace. Understanding the best balance of performance and compression is an important factor to consider when designing a mixed compressed environment.

RTC can be purchased to run on one control enclosure (I/O Group) of a cluster and not for another allowing for shared environments to exist depending on need and providing environments that meet the needs.

To gain insight of the what should be expected a tool has been created which can be used to gather details and provide input on the workload type and pattern to show its level of compressibility. An IBM tool called the Comprestimator can be used to analyze data gathered.

For more information about the usage and capabilities of this feature, see *Real-time Compression in SAN Volume Controller and Storwize V7000*, REDP-4859, at the following website:

<http://w3.itso.ibm.com/redpieces/abstracts/redp4859.html?open>

2.5.14 iSCSI

iSCSI is an alternative means of attaching hosts to the IBM Flex System V7000 Storage Node. All communications with external back-end storage subsystems, or other IBM virtual storage systems must be done through a Fibre Channel or FCOE connection.

The iSCSI function is a *software function* that is provided by the IBM Flex System V7000 Storage Node code, not hardware.

In the simplest terms, iSCSI allows the transport of SCSI commands and data over a TCP/IP network, based on IP routers and Ethernet switches. iSCSI is a block-level protocol that encapsulates SCSI commands into TCP/IP packets and uses an existing IP network, instead of requiring expensive FC HBAs and a SAN fabric infrastructure.

A pure SCSI architecture is based on the client/server model. A client (for example, server or workstation) initiates read or write requests for data from a target server (for example, a data storage system).

Commands, which are sent by the client and processed by the server, are put into the Command Descriptor Block (CDB). The server runs a command, and completion is indicated by a special signal alert.

The major functions of iSCSI include encapsulation and the *reliable delivery* of CDB transactions between initiators and targets through the Internet Protocol network, especially over a potentially unreliable IP network.

The concepts of names and addresses have been carefully separated in iSCSI:

- ▶ An *iSCSI name* is a location-independent, permanent identifier for an iSCSI node. An iSCSI node has one iSCSI name, which stays constant for the life of the node. The terms *initiator name* and *target name* also refer to an iSCSI name.
- ▶ An *iSCSI address* specifies not only the iSCSI name of an iSCSI node, but also a location of that node. The address consists of a host name or IP address, a TCP port number (for the target), and the iSCSI name of the node. An iSCSI node can have any number of addresses, which can change at any time, particularly if they are assigned by way of Dynamic Host Configuration Protocol (DHCP). An IBM Flex System V7000 Storage Node control canister represents an iSCSI node and provides statically allocated IP addresses.

Each iSCSI node, that is, an initiator or target, has a unique iSCSI Qualified Name (IQN), which can have a size of up to 255 bytes. The IQN is formed according to the rules adopted for Internet nodes.

The iSCSI qualified name format defined in RFC3720 contains these elements (in order):

- ▶ The string “iqn”.
- ▶ A date code specifying the year and month in which the organization registered the domain or subdomain name used as the naming authority string.
- ▶ The organizational naming authority string, which consists of a valid, reversed domain or a subdomain name.
- ▶ Optionally, a colon (:), followed by a string of the assigning organization’s choosing, which must make each assigned iSCSI name unique.

For IBM Flex System V7000 Storage Node the IQN for its iSCSI target is specified as:

```
iqn.1986-03.com.ibm:2145.<clustername>.<nodename>
```

On a Windows server, the IQN, that is, the name for the iSCSI initiator, can be defined as:

```
iqn.1991-05.com.microsoft:<computer name>
```

The IQNs can be abbreviated by using a descriptive name, known as an *alias*. An alias can be assigned to an initiator or a target. The alias is independent of the name and does not need to be unique. Because it is not unique, the alias must be used in a purely informational way. It cannot be used to specify a target at login or used during authentication. Both targets and initiators can have aliases.

An iSCSI name provides the correct identification of an iSCSI device irrespective of its physical location. Remember that the IQN is an *identifier*, not an *address*.

Changing names: Observe the following precautions when changing system or node names for an IBM Flex System V7000 Storage Node clustered system that has servers connected to it using SCSI. Be aware that because the system and node name are part of the IQN for the IBM Flex System V7000 Storage Node, you can lose access to your data by changing these names. The IBM Flex System V7000 Storage Node GUI shows a specific warning, but the CLI does *not*.

The iSCSI session, which consists of a *login phase* and a *full feature phase*, is completed with a special command.

The login phase of the iSCSI is identical to the FC port login process (PLOGI). It is used to adjust various parameters between two network entities and to confirm the access rights of an initiator.

If the iSCSI login phase is completed successfully, the target confirms the login for the initiator; otherwise, the login is not confirmed and the TCP connection breaks.

As soon as the login is confirmed, the iSCSI session enters the full feature phase. If more than one TCP connection was established, iSCSI requires that each command / response pair goes through one TCP connection. Thus, each separate read or write command is carried out without the necessity to trace each request for passing separate flows. However, separate transactions can be delivered through separate TCP connections within one session.

For further details about configuring iSCSI, see Chapter 11, “SAN connections and configuration” on page 431.

2.5.15 Fibre Channel over Ethernet (FCoE)

Fibre Channel over Ethernet (FCoE) is a standard specified by ANSI T11 committee within the FC-BB-5 for enabling the transmission of FC protocol and data across an Ethernet network. As shown in Figure 2-18, the new environment supports both FCoE and TCP/IP traffic to be able to share a common Ethernet network.

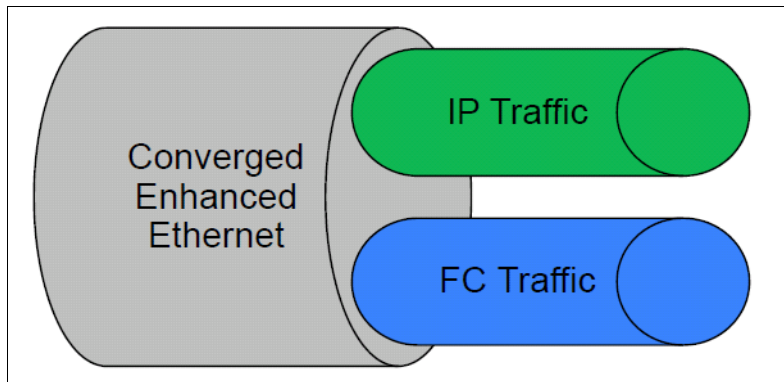


Figure 2-18 New enhanced Ethernet environment support

Table 2-7 shows the differences between the use of FCoE and iSCSI for transfer environments.

Table 2-7 FCoE and iSCSI differences

FCoE	iSCSI
Local-area, lossless links, no routing allowed	Allows many hops, lossy connections and high latency
Simple encapsulation of Fibre Channel	Substantial complexity on top of TCP
Low overhead – similar to Fibre Channel	Overhead varies - typically higher
Storage administrators know Fibre Channel well	TCP/iP well understood by most IT staff
Frame loss can quickly become catastrophic	Frame loss recovery built into protocol stack

FCoE allows for fewer network adapters to be required, as both protocols can share the same adapter, reducing hardware cost and freeing up bus slots in hosts.

For in-depth details and recommendations for using these protocols, see *Storage and Network Convergence Using FCoE and iSCSI*, SG24-7986 available at this website:

<http://w3.itso.ibm.com/abstracts/sg247986.html?0pen>

2.6 Advanced copy services

IBM Flex System V7000 Storage Node supports the following copy services:

- ▶ FlashCopy
- ▶ Synchronous Remote Copy
- ▶ Asynchronous Remote Copy

2.6.1 FlashCopy

FlashCopy makes a copy of a source volume on a target volume. The original content of the target volume is lost. After the copy operation has started, the target volume has the contents of the source volume as they existed at a single point in time. Although the copy operation takes time, the resulting data at the target appears as though the copy was made instantaneously.

FlashCopy is sometimes described as an instance of a time-zero (T0) copy or a point in time (PiT) copy technology.

FlashCopy can be performed on multiple source and target volumes. FlashCopy permits the management operations to be coordinated so that a common single point in time is chosen for copying target volumes from their respective source volumes.

IBM Flex System V7000 Storage Node also permits multiple target volumes to be FlashCopied from the same source volume. This capability can be used to create images from separate points in time for the source volume, and to create multiple images from a source volume at a common point in time. Source and target volumes can be thin-provisioned volumes.

Reverse FlashCopy enables target volumes to become restore points for the source volume without breaking the FlashCopy relationship and without waiting for the original copy operation to complete. IBM Flex System V7000 Storage Node supports multiple targets and thus multiple rollback points.

Most clients aim to integrate the FlashCopy feature for point in time copies and quick recovery of their applications and databases. IBM Support is provided by Tivoli Storage FlashCopy Manager, which is described at the following website:

<http://www-01.ibm.com/software/tivoli/products/storage-flashcopy-mgr/>

You can read a detailed description about the FlashCopy copy services in 9.2, “FlashCopy” on page 342.

2.6.2 IBM Flex System V7000 Remote Mirroring software

The IBM Flex System V7000 Remote Mirroring software provides both Metro and Global mirroring capabilities between the IBM Flex System V7000 Storage Nodes, or between IBM Flex System V7000 Storage Nodes and IBM Storwize V7000s or IBM SAN Volume Controllers (SVC). This capability means that customers have greater flexibility in their expanding environments using Flex System V7000, Storwize V7000, and SVC with the ability now to remote mirror from one system to the other. With the wide variety of storage systems that can be managed under any of these systems, the options available for replicating data between a number of storage systems is multiplied. Remote deployments for disaster recovery for current SVC and Storwize V7000 environments can easily be fitted with the Flex System V7000 or vice versa.

The Copy Services layer sits above and operates independently of the function or characteristics of the underlying disk subsystems used to provide the storage resources to the IBM Flex System V7000 Storage Node. Figure 2-19 shows an example of possible copy service relationships (they must be at Version 6.4 or later).

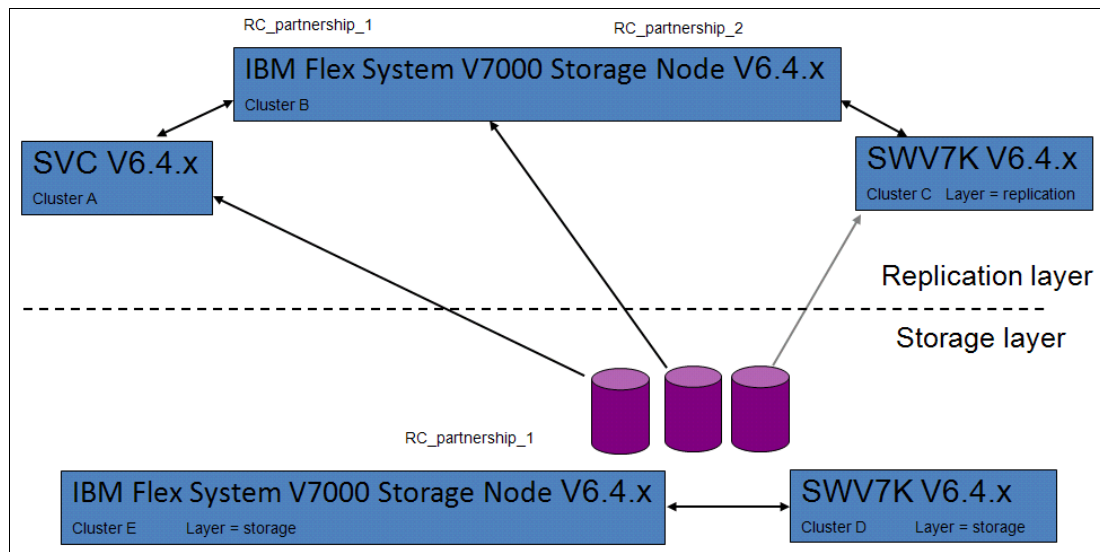


Figure 2-19 Example of possible copy services relationships

With Metro Mirroring, IBM Flex System V7000 Remote Mirroring provides synchronous data replication at distances less than 300 Km. This capability is supported across either an FC or FCoE SAN fabric infrastructure. Mirroring over iSCSI connections is not currently supported.

Global Mirroring provides for long distances of greater than 300 Km and is only supported over an FC SAN infrastructure. Tunneling over a WAN network is frequently used for the greater distance needs.

Customers who want to use the global mirror capability with Flex System V7000 on a low bandwidth link between sites can do so with the use of the low bandwidth remote mirroring. This capability provides options to help administrators balance network bandwidth requirements and RPO times for applications, helping reduce operation costs for disaster recovery solutions. Remote mirroring supports higher RPO times by allowing the data at the disaster recovery site to get further out of sync with the production site if the communication link limits replication, and then approaches synchronicity again when the link is not as busy. This low bandwidth remote mirroring uses space-efficient FlashCopy targets as sources in remote copy relationships to increase the time allowed to complete a remote copy data cycle.

See 9.3, “Remote Copy” on page 380 for more information about the Remote Copy services.

2.6.3 Remote Copy licensing

To authorize use of remote copy capabilities of the IBM Flex System V7000, you must purchase a feature license for IBM Flex System V7000 Remote Mirroring Software. This license is based on the quantity of enclosures that are authorized for the feature.

To determine the number of authorized enclosures needed, count all enclosures on the system that are managed by the IBM Flex System V7000 Disk System, including each internal enclosure licensed with the IBM Flex System V7000 Base Software and each external enclosure licensed with the IBM Flex System V7000 External Virtualization Software. Counting should be done at both the ends of the mirroring relationship. The location with the least number of members is the location that defines the number of licenses needed for each location site.

For example, if you have site A with a DS5020 consisting of three enclosures managed by an IBM Flex System V7000 consisting of one control enclosure and one expansion enclosure mirroring to a partner at site B with an IBM Flex System V7000 consisting of one control enclosure and two expansion enclosures and no external virtual storage system attached, then in order to authorize remote mirroring for this configuration you will need a license to the IBM Flex System V7000 Remote Mirroring Software, for a quantity of six enclosures; three for site A and three for site B. If however, in the future you were to grow your site B system to include more enclosures or an external virtualized system you would then need to re-evaluate the licensing to add the appropriate number for the new configuration of the site with the fewest enclosures managed.

For a remote mirroring Flex System V7000 with an SVC in the mix. because of a difference in metrics, SVC mirroring can be licensed for a subset of the total storage virtualized. However, the IBM Flex System V7000 and the IBM Storwize V7000 still must use the aforementioned model for their licensing.

2.6.4 Synchronous / Asynchronous Remote Copy

The general application of Remote Copy seeks to maintain two copies of data. Often, the two copies are separated by distance, but not always.

The Remote Copy can be maintained in one of two modes: synchronous or asynchronous.

With the IBM Flex System V7000 Storage Node, Metro Mirror and Global Mirror are the IBM branded terms for the functions that are synchronous Remote Copy and asynchronous Remote Copy.

Synchronous Remote Copy ensures that updates are committed at both the primary and the secondary before the application considers the updates complete; therefore, the secondary is fully up to date if it is needed in a failover. However, the application is fully exposed to the latency and bandwidth limitations of the communication link to the secondary. In a truly remote situation, this extra latency can have a significant adverse effect on application performance.

Special configuration guidelines exist for SAN fabrics that are used for data replication. It is necessary to consider the distance and available bandwidth total network links to determine the appropriate method to use. See 9.3, "Remote Copy" on page 380 for details on planning, configuring and using Remote Copy for replication functions.

With the Global Mirror method, there is a design option that assists with low bandwidth for the IBM Flex System V7000 Storage Node and the other IBM virtual storage systems that are mirroring to it running 6.4.1 code level. This option uses change volumes associated with the primary and secondary volumes. These volumes are used to record changes to the remote copy volume, the FlashCopy relationship that exists between the secondary volume and the change volume, and between the primary volume and the change volume. This function is called Global Mirror cycling mode. Figure 2-20 shows an example of this function where you can see the relationship between volumes and change volumes.

Support: Though this feature was added in the 6.3 code level, only 6.4.1 code level is supported to mirror to the IBM Flex System V7000 Storage Node.

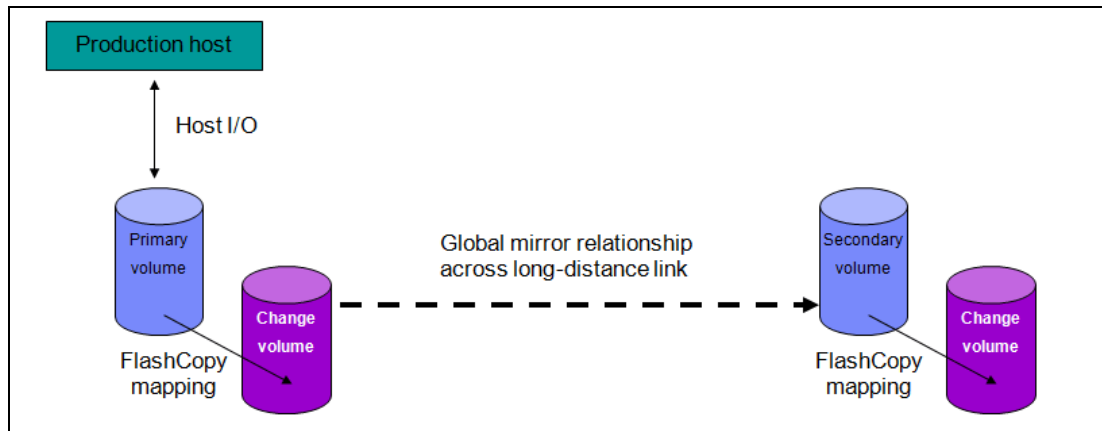


Figure 2-20 Global Mirror with change volumes

In asynchronous Remote Copy, the application acknowledges that the write is complete before the write is committed at the secondary. Hence, on a failover, certain updates (data) might be missing at the secondary. The application must have an external mechanism for recovering the missing updates, if possible. This mechanism can involve user intervention. Recovery on the secondary site involves starting the application on this recent “backup” and then rolling forward or backward to the most recent commit point.

2.6.5 Copy Services configuration limits

In Table 2-8, we describe the Copy Services configuration limits. For the most up-to-date list of these limits, see the following website:

<http://www.ibm.com/support/docview.wss?uid=ssg1S1003702&myns=s028&mynp=familyind5402112&mync=E>

Table 2-8 Copy Services configuration limits

Properties	Maximum number	Comments
Remote Copy (Metro Mirror and Global Mirror) relationships per clustered system	4096	This configuration can be any mix of Metro Mirror and Global Mirror relationships.
Remote Copy relationships per consistency group	4096	No additional limit is imposed beyond the Remote Copy relationships per clustered system limit.

Properties	Maximum number	Comments
Remote Copy consistency groups per clustered system	256	-
Total Metro Mirror and Global Mirror volume capacity per I/O group	1024 TB	This limit is the total capacity for all master and auxiliary volumes in the I/O group.
FlashCopy mappings per clustered system	4096	-
FlashCopy targets per source	256	-
Cascaded Incremental FlashCopy maps	4	A volume can be the source of up to four incremental FlashCopy maps. If this number of maps is exceeded, then the FlashCopy behavior for that cascade becomes non-incremental.
FlashCopy mappings per consistency group	512	-
FlashCopy consistency groups per clustered system	127	-
Total FlashCopy volume capacity per I/O group	1024 TB	4096 for a full four node clustered system with four I/O groups.

2.7 Management and support tools

The IBM Flex System V7000 Storage Node system can be managed through the IBM Flex System Management Node (FSM), or by using the native management software that runs in the hardware itself.

The FSM simplifies storage management in the following ways:

- ▶ Centralizes the management of storage network resources with IBM storage management software.
- ▶ Provides greater synergy between storage management software and IBM storage devices.
- ▶ Reduces the number of servers that are required to manage your software infrastructure.
- ▶ Provides higher-level functions.

2.7.1 IBM Assist On-site and remote service

The IBM Assist On-site tool is a remote desktop-sharing solution that is offered through the IBM website. With it, the IBM service representative can remotely view your system to troubleshoot a problem.

You can maintain a chat session with the IBM service representative so that you can monitor this activity and either understand how to fix the problem yourself or allow the representative to fix it for you.

To use the IBM Assist On-site tool, the SSPC or master console must be able to access the Internet. The following website provides further information about this tool:

<http://www.ibm.com/support/assistsite/>

When you access the website, you sign in and enter a code that the IBM service representative provides to you. This code is unique to each IBM Assist On-site session. A plug-in is downloaded on to your SSPC or master console to connect you and your IBM service representative to the remote service session. The IBM Assist On-site contains several layers of security to protect your applications and your computers.

You can also use security features to restrict access by the IBM service representative.

Your IBM service representative can provide you with more detailed instructions for using the tool.

2.7.2 Event notifications

IBM Flex System V7000 Storage Node can use Simple Network Management Protocol (SNMP) traps, syslog messages, and a Call Home email to notify you and the IBM Support Center when significant events are detected. Any combination of these notification methods can be used simultaneously.

Each event that IBM Flex System V7000 Storage Node detects is assigned a notification type of Error, Warning, or Information. You can configure IBM Flex System V7000 Storage Node to send each type of notification to specific recipients.

2.7.3 SNMP traps

SNMP is a standard protocol for managing networks and exchanging messages. IBM Flex System V7000 Storage Node can send SNMP messages that notify personnel about an event. You can use an SNMP manager to view the SNMP messages that IBM Flex System V7000 Storage Node sends. You can use the management GUI or the IBM Flex System V7000 Storage Node command-line interface to configure and modify your SNMP settings.

You can use the Management Information Base (MIB) file for SNMP to configure a network management program to receive SNMP messages that are sent by the IBM Flex System V7000 Storage Node. This file can be used with SNMP messages from all versions of IBM Flex System V7000 Storage Node Software.

2.7.4 Syslog messages

The syslog protocol is a standard protocol for forwarding log messages from a sender to a receiver on an IP network. The IP network can be either IPv4 or IPv6. IBM Flex System V7000 Storage Node can send syslog messages that notify personnel about an event. IBM Flex System V7000 Storage Node can transmit syslog messages in either expanded or concise format. You can use a syslog manager to view the syslog messages that IBM Flex System V7000 Storage Node sends. IBM Flex System V7000 Storage Node uses the User Datagram Protocol (UDP) to transmit the syslog message. You can use the management GUI or the IBM Flex System V7000 Storage Node command-line interface to configure and modify your syslog settings.

2.7.5 Call Home email

The Call Home feature transmits operational and error-related data to you and IBM through a Simple Mail Transfer Protocol (SMTP) server connection in the form of an event notification email. When configured, this function alerts IBM service personnel about hardware failures and potentially serious configuration or environmental issues. You can use the Call Home function if you have a maintenance contract with IBM or if the IBM Flex System V7000 Storage Node is within the warranty period.

To send email, you must configure at least one SMTP server. You can specify as many as five additional SMTP servers for backup purposes. The SMTP server must accept the relaying of email from the IBM Flex System V7000 Storage Node clustered system IP address. You can then use the management GUI or the command-line interface to configure the email settings, including contact information and email recipients. Set the reply address to a valid email address. Send a test email to check that all connections and infrastructure are set up correctly. You can disable the Call Home function at any time using the management GUI or the command-line interface as well.

2.8 Useful references from Storwize V7000 websites

Due to the common design of the software and procedures of the IBM Storwize V7000 and the IBM Flex System V7000 Storage Node, the following information is provided for additional references.

- ▶ IBM Flex System Information Center:
<http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp>
- ▶ IBM Flex System V7000 Storage Node Information Center:
http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.4939.doc/site_product_page.html
- ▶ IBM Flex System V7000 Storage Node - Learning Modules:
- ▶ http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.4939.learning/site_learning_tutorials.html
- ▶ IBM Flex System V7000 Storage Node Installation Guide (GC27-4203):
http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.4939.doc/site_bkmap_quickinstbk.pdf
- ▶ IBM Storwize V7000 Expansion Enclosure Installation Guide, Machine type 2076 (GC27-4234):
http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.4939.doc/tbrd_bkmap_install_2076.pdf
- ▶ IBM Flex System V7000 Storage Node Troubleshooting, Recovery, and Maintenance Guide (GC27-4205):
http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.4939.doc/site_bkmap_trblshoot.pdf
- ▶ IBM Flex System V7000 Storage Node Support home:
http://www.ibm.com/support/entry/portal/overview/hardware/puresystems/pureflex_system/storage_node/flex_system_v7000

- ▶ The IBM Storwize V7000 Storage Support page is at the following website:
<http://www.ibm.com/storage/support/storwize/v7000>
- ▶ The IBM Storwize V7000 Storage home page is at the following website:
http://www-03.ibm.com/systems/storage/news/center/storwize_v7000/index.html
- ▶ The Storwize V7000 Storage Supported hardware list is at the following website:
<http://www-01.ibm.com/support/docview.wss?uid=ssg1S1003908>
- ▶ The Storwize V7000 Storage Configuration Limit and Restrictions are at the following website:
<http://www-01.ibm.com/support/docview.wss?uid=ssg1S1003902>
- ▶ The Storwize V7000 Storage Online documentation is at the following website:
http://www-947.ibm.com/support/entry/portal/Documentation/Hardware/System_Storage/Disk_systems/Mid-range_disk_systems/IBM_Storwize_V7000_%282076%29
- ▶ The Storwize V7000 Storage Online Information Center is at the following website:
<http://publib.boulder.ibm.com/infocenter/storwize/ic/index.jsp>
- ▶ You can see IBM Redbooks publications about both the IBM Flex System and IBM Storwize V7000 Storage at the following website:
<http://www.redbooks.ibm.com/cgi-bin/searchsite.cgi?query=v7000>

2.9 IBM virtual storage learning videos on YouTube

Some interesting videos are available on YouTube that describe the IBM Storwize V7000 Storage. With the common software design shared with this system and the IBM Flex System V7000 Storage Node many of these can be used to help develop a familiarity with the terminology and many of the GUI and management processes. They are available at the websites shown in Table 2-9.

Table 2-9 Videos available on YouTube

Description of the video	Link WEB
IBM Storwize V7000 Storage Virtualization Terminology Overview	http://www.youtube.com/watch?v=I2rzt3m2gP0
IBM Storwize V7000 Interface tour	http://www.youtube.com/watch?v=FPbNRs9HacQ
IBM Storwize V7000 Volume management	http://www.youtube.com/watch?v=YXeKqH8Sd9o
IBM Storwize V7000 Migration	http://www.youtube.com/watch?v=dXxnUN6dk74
IBM Storwize V7000 Introduction to FlashCopy	http://www.youtube.com/watch?v=MXWgGwjBzG4
IBM Storwize V7000 Introduction to Remote Copy 1	http://www.youtube.com/watch?v=koFYm--gnEc
IBM Storwize V7000 Introduction to Remote Copy 2	http://www.youtube.com/watch?v=kM0FFS3NYas
IBM Storwize V7000 Rapid Application Storage Solution - Tivoli Storage FlashCopy Manager	http://www.youtube.com/watch?v=1zgEljRVDSc
VMware data protection with Storwize V7000	http://www.youtube.com/watch?v=vec0ap-qwbA
IBM SAN Volume Controller and Storwize V7000 Performance Panel Sped-up! (HD)	http://www.youtube.com/watch?v=7noC71tLkWs



Systems management

This chapter provides an overview of storage management for the IBM Flex System V7000 Storage Node from the IBM Flex System Chassis Management Module and the IBM Flex System Manager Node, and shows how to use the navigation tools.

While there are several ways to manage the various components of the IBM Flex Systems, a description of all these options is beyond the scope of this publication. In this chapter, we cover the management options related to IBM Flex System V7000 Storage Node and the I/O modules. For detailed information regarding other management options, see *Implementing Systems Management of IBM PureFlex System*, SG24-8060.

This chapter describes the following topics:

- ▶ 3.1, “System Management overview” on page 82
- ▶ 3.2, “IBM Flex System Chassis Management Module (CMM)” on page 88
- ▶ 3.3, “Flex System Manager (FSM)” on page 110

3.1 System Management overview

The IBM PureFlex System is designed to improve efficiency in system management. A variety of functions and features are designed towards providing a simplified yet dynamic management experience to the user:

- ▶ Advanced automation of configuration and management
- ▶ Enhanced and user-friendly interface, cross-resource integration
- ▶ Easier monitoring and problem management
- ▶ Policy-based expert patterns

3.1.1 Integrated platform management tools

System Management can be performed by using easy-to-use integrated platform management tools (CMM and IMMv2), as well as other popular system management software (IBM Tivoli and IBM Director), and third-party enterprise management tools (CA, BMC, and HP). Due to the flexibility of an IBM Flex System, many management features and functions such as security, monitoring, provisioning, and automation, can be plugged easily into an existing IT infrastructure.

Within an IBM Flex System Enterprise Chassis, each node has a management controller that communicates by the CMM-enabled 1 GbE private management network allowing centralized management capability. Different chassis components supported in the IBM Flex System Enterprise Chassis implement different management controllers at node-level as well.

The following set of hardware management devices and interfaces are commonly used for managing an IBM Flex System:

- ▶ IBM Flex System Manager (FSM)
- ▶ IBM Flex System Chassis Management Module (CMM)
- ▶ Switch modules (IO modules)
- ▶ Compute node management processors
 - Integrated Management Module (IMMv2) for intel-based servers
 - Flexible Service Processor (FSP) for POWER-based servers

IBM Flex System Manager (FSM)

FSM is a highly integrated management device that offers single-system management across physical and virtual resources. It provides complete control of IBM Flex System components and features. It offers the following capabilities:

- ▶ Virtualization management
- ▶ Resource allocation and control
- ▶ Network and storage control

More information regarding IBM Flex System Manager can be found at 3.3.1, “Overview of IBM Flex System Manager (FSM)” on page 110.

IBM Flex System Chassis Management Module (CMM)

CMM is the default integrated platform for managing a single IBM Flex System Enterprise Chassis. It occupies one of the two chassis management bays on the rear of the chassis. A second module can be installed for redundancy.

More information regarding IBM Flex System Chassis Management Module can be found at 3.2.1, “Overview of IBM Flex System Chassis Management Module” on page 88.

I/O modules (switch modules)

I/O modules contain the internal management interfaces and Inter-Integrated Circuit (I2C) for management. They can be managed by the IBM Flex System Manager, IBM Flex System Chassis Management Module, command line interface (CLI), web interface, IBM System Network Element Manager (SNEM), or any third-party SNMP-based management tool.

Compute node management processors

There are two types of compute nodes in an IBM Flex System:

- ▶ For IBM x86 (intel-based) compute nodes, the management interface is IMMv2, which is a powerful management interface providing control at node-level.
- ▶ For POWER compute node, Flexible Service Processor(FSP) provides a serial-over-LAN (SOL) interface for performing common management tasks on the node.

Table 3-1 gives an overview of system management capabilities of these hardware management devices.

Table 3-1 System management

IBM Flex System Manager	IBM Flex System Chassis Management Module	IMMv2 / Flexible Service Processor	Network Switch
<ul style="list-style-type: none"> ▶ Remote access by browser ▶ Multi-Chassis Rack Management ▶ Configuration and Updates ▶ Network Configuration Manger ▶ Service & Support Manager ▶ Virtualization Management (VM Control) ▶ Active Energy Manager ▶ Local Console for Chassis Management ▶ Remote control for Compute node(s) 	<ul style="list-style-type: none"> ▶ Aggregation point for chassis ▶ Integrated Remote Chassis Management ▶ Power and Thermal Management ▶ Internal Chassis Network with Management packet forwarding ▶ Web, CLI, SNMP, CIM-XML ▶ CMM redundancy (optional additional CMM) ▶ RAS, FFDC, Service &Support, Chassis Event Logs 	<ul style="list-style-type: none"> ▶ Remote Presence (CLI/video/USB) ▶ Power and Thermal Management. ▶ Inventory, Fault detection ▶ Update Management ▶ Configuration Management ▶ RAS (reliability, availability, serviceability), FFDC (First-failure data capture), Logs 	<ul style="list-style-type: none"> ▶ Power, Thermal sensing ▶ Update Management ▶ Inventory, Fault detection ▶ Configuration Management ▶ RAS, FFDC (First-failure data capture), Logs ▶ SNMP, CLI, Web, other proprietary protocols

3.1.2 IBM Flex System storage management

The IBM Flex System offers strong storage management capabilities that allow you to gain advanced functionality with storage nodes in your system while also taking advantage of your existing storage infrastructure through advanced virtualization.

Storage management features

An IBM Flex System includes these advanced storage management features:

- ▶ Automated discovery and inventory: Physical seating of Storage Node triggers discovery and inventory, automatically adding it to IBM Flex System Manager
- ▶ Intelligent storage provisioning: Integrated storage allocation
- ▶ Automated firmware update: Update Manager provides firmware level compliance monitoring and management
- ▶ Integrated event monitoring: Seamless visibility to events and status
- ▶ Automatic credential exchange: Storage Node's initial credentials are set by IBM Flex System Manager

IBM Flex System V7000 Storage Node

The IBM Flex System V7000 Storage Node is physically and logically integrated into an IBM Flex System. IBM Flex System V7000 Storage Node is designed on the industry-leading storage virtualization and efficiency capabilities of IBM Storwize V7000 by simplifying and speeding deployment, and providing greater integration of server and storage management to automate and streamline provisioning for greater responsiveness to business needs and lower overall cost.

IBM Flex System V7000 Storage Node is an Integrated storage system with advanced storage efficiency capabilities, which include Thin Provisioning, FlashCopy, Easy Tier, Real-time Compression, and Non-disruptive Migration. It offers external virtualization for rapid data center integration, Metro and Global Mirror for multi-site recovery, scalability up to 240 small form-factor (SFF) drives (HDD and SSD), and clustered systems support up to 960 drives.

IBM Flex System V7000 Storage Node simplifies storage administration with a single user interface for all your storage with a management console that is integrated with the comprehensive management system. The enhanced management capabilities allow you to virtualize third-party storage with non-disruptive migration of the existing storage infrastructure. Intelligent tiering provides the advantage of balancing high performance and cost for your storage needs. It also protects data for business continuity and disaster recovery with built-in local and remote replication and snapshot functions.

IBM Flex System V7000 Storage Node enhancements include physical chassis plug-and-play integration, automated deployment and discovery, integration into Flex System Manager Chassis Map, and FCoE optimized offering in addition to FC and iSCSI.

For storage management, the IBM Flex System V7000 Storage Node system can have two management IP addresses:

- ▶ One for the 1-Gb/s internal management network (required)
- ▶ One on the first 10-Gb/s Ethernet port in the configuration node canister (optional)

3.1.3 Storage management interfaces

The IBM Flex System storage infrastructure contains a variety of management interfaces that can be classified into four distinct levels based on their level of depth and detail in management capabilities:

Starting Level: IBM Flex System Manager Chassis Map

The IBM Flex System Manager offers a Chassis Map, which is the single point of management and starting point for IBM Flex System. It allows for physical navigation through the entire system, as shown in Figure 3-1.



Figure 3-1 FSM Chassis Map

Upper Level: IBM Flex System Manager Storage Control

Higher level datacenter management can be done using IBM Flex System Manager Storage Control feature. It provides integrated management with a Systems and Storage perspective. Using Storage Control, you can manage storage across your data center connected to the IBM Flex System, and also map connections between servers, switches and storage. Storage Control is shown in Figure 3-2.



Figure 3-2 IBM Flex System Manager Storage control

Detailed Level: IBM Flex System V7000 Storage Node User Interface

This interface is inherent in every IBM Flex System V7000 Storage Node. It allows for detailed disk setup (RAID arrays and LUNs). All the advanced disk features like Easy Tier, Mirroring, Clustering, and Virtualization can be configured using the IBM Flex System V7000 Storage Node User Interface. Figure 3-3 shows the IBM Flex System V7000 Storage Node interface.

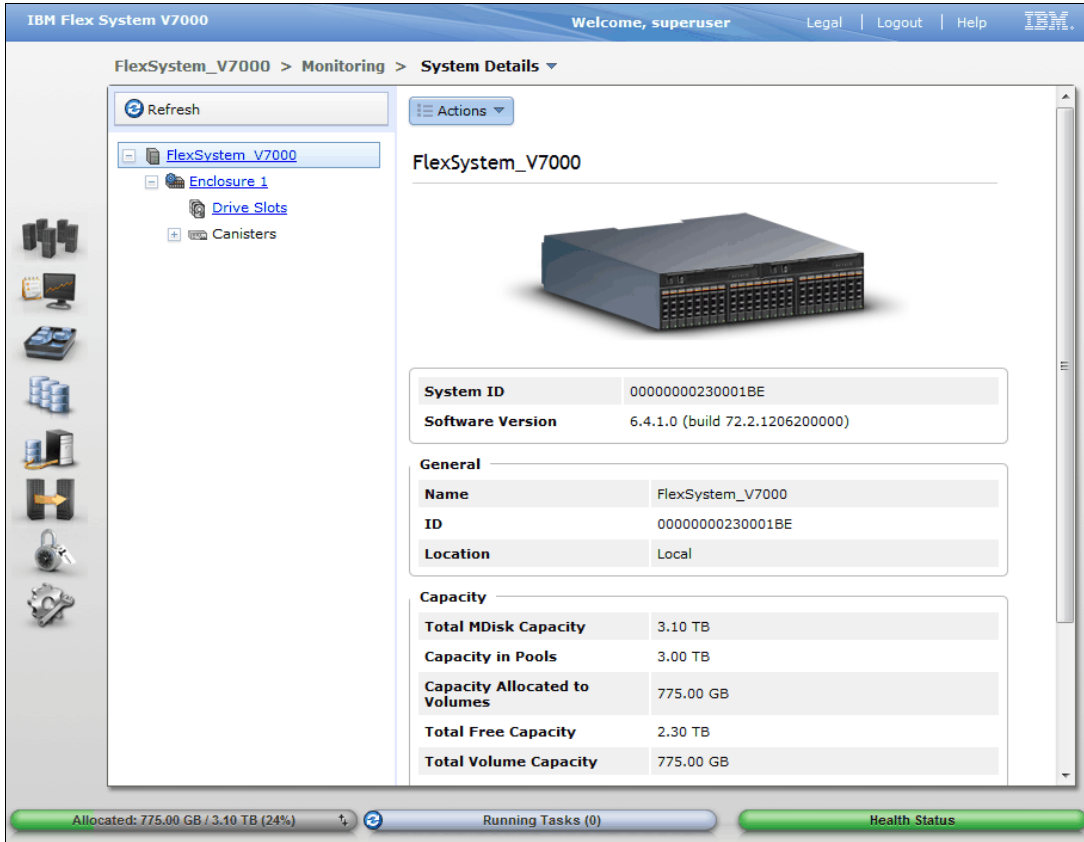


Figure 3-3 IBM Flex System V7000 GUI - System Details

For further details on management of an IBM Flex System V7000 Storage Node, see Chapter 5, “IBM Flex System V7000 Storage Node GUI interface” on page 175.

3.2 IBM Flex System Chassis Management Module (CMM)

The IBM Flex System Chassis Management Module (CMM) is a hot-swap module that configures and manages all installed chassis components. The chassis comes with one IBM Flex System Chassis Management Module in the CMM bays.

The IBM Flex System Chassis Management Module auto-detects and manages the hardware of the entire IBM Flex System Enterprise Chassis. It communicates with the management controller in each node (IMMv2 in intel processor-based computer nodes, FSP in power Power7 processor-based nodes), along with the management applications of storage node and I/O modules to provide centralized system monitoring and management.

3.2.1 Overview of IBM Flex System Chassis Management Module

A IBM Flex System Chassis Management Module (CMM) gives control over the IBM Flex System at the chassis level, simplifying installation and management of your installation.

The CMM is central to the management of the chassis, and is required in the IBM Flex System Enterprise Chassis. One CMM comes standard with the chassis. There are two CMM bays in the chassis, giving the user an option to add a second CMM for redundancy.

Through an embedded firmware stack, the CMM implements functions to monitor, control, and provide external user interfaces to manage all chassis resources. Here are the key functions of the IBM Flex System Chassis Management Module:

- ▶ Single-chassis management
- ▶ Initial setup configuration wizard from http/s interface
- ▶ Detailed system health status for components, along with suggested recovery actions
- ▶ Chassis level hardware inventory management
- ▶ Chassis configuration and validation
- ▶ System and component status information
- ▶ Component LEDs states and controls
- ▶ Event handling and notification
- ▶ Hardware access states
- ▶ Fan status and control
- ▶ Thermal management and monitoring
- ▶ Power management and control
- ▶ Switch modules - monitor and control access
- ▶ Security features and policies
- ▶ Service data collection and Call Home services
- ▶ Configuration backup and restore
- ▶ Firmware updates
- ▶ Low level service processor resets
- ▶ Nodes and I/O module configuration and management
- ▶ Diagnostics
- ▶ Network management
- ▶ Chassis and compute node initialization

Tip: The IBM Flex System Enterprise Chassis supports up to 84 local CMM user accounts, and allows up to 32 simultaneous sessions.

Management interfaces

The IBM Flex System Chassis Management Module allows two user interfaces to perform management functions:

- ▶ The IBM Flex System Chassis Management Module web-based graphical user interface (GUI) provides a way to perform CMM functions within a supported web browser. This user interface is a program that is part of the firmware on IBM Flex System Chassis Management Module.
- ▶ The IBM Flex System Chassis Management Module command-line interface (CLI) provides direct access to IBM Flex System management functions as an alternative to using the web-based user interface.

The CMM CLI (SSH/Telnet) Web and SNMP management interface is accessible remotely over a Ethernet network. No crossover cable is required to connect to CMM if a local PC is used, as CMM performs automatic media dependent interface crossover.

Web-based Graphical User Interface (GUI)

The IBM Flex System Chassis Management Module web interface communicates with the management program to perform chassis management tasks:

- ▶ Controlling the chassis, compute nodes, and other components
- ▶ Defining login IDs and passwords
- ▶ Discovering other chassis on the network and enabling access to them
- ▶ Accessing and configuring I/O modules
- ▶ Monitoring the status of the compute nodes and other components
- ▶ Selecting recipients for alert notification of specific events
- ▶ Using a remote console for the compute nodes
- ▶ Changing the startup sequence in a compute node
- ▶ Configuring security settings such as data encryption and user account security

For detailed overview of the IBM Flex System Chassis Management Module web-based user interface, visit the following website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.cmm.doc/cmm_user_guide.html

Command-line interface (CLI)

Using the IBM Flex System Chassis Management Module command-line interface (CLI), you can issue commands to control the power and configuration of the CMM and other components that are in an IBM Flex System Enterprise Chassis.

The command-line interface also provides access to the text-console command prompt on each compute node through a Serial Over LAN (SOL) connection.

Access: You can access the CMM CLI through a direct serial or Ethernet connection to the CMM, through a Telnet connection to the IP address of the CMM, or through a Secure Shell (SSH) connection to the CMM.

You can find detailed reference information about using the command-line interface at the following website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.cmm.doc/cmm_cli_guide.htm

3.2.2 Accessing the CMM

After initial connection, you need to configure the IBM Flex System Enterprise Chassis and its components for your operating environment. When the chassis is started, all installed components will be discovered and the vital product data (VPD) of each component is stored in the IBM Flex System Chassis Management Module.

Introduction to the CMM

The CMM will automatically configure the remote management port of the CMM so that you can configure and manage all of the chassis components. By default, the CMM is configured to respond to Dynamic Host Configuration Protocol (DHCP) first before using its static IPv4 address. It has the following default settings:

IP address 192.168.70.100
Subnet 255.255.255.0
User ID USERID (all capital letters)
Password PASSWORD (all capital letters, with a zero instead of the letter O)

Tip: The client computer that you connect to the CMM must be configured to operate on the same subnet as the CMM. The IP address of the CMM must also be in the same local domain as the client computer.

The default information, MAC address, and IPv6 link-local address is available in the network access card attached to all new CMMs, as shown in Figure 3-4.




Figure 3-4 Network Access Tag

Tip: The default security setting is *secure*, so HTTPS is required to connect to the CMM.

Logging in to the CMM

The following section gives step-by-step instructions regarding logging in to the CMM.

1. Open a web browser, and type in the IP address of the CMM.
2. At the Login panel, enter your User name and Password, and set the Inactive session time-out value, as shown in Figure 3-5.

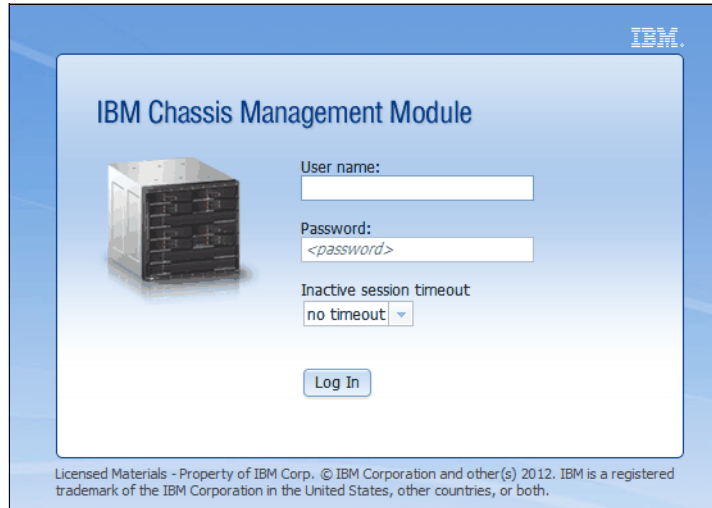


Figure 3-5 IBM Chassis Management Module Login panel

An Initial Setup wizard starts automatically the first time you access the web interface of a new CMM or a CMM that has been reset to its default settings. For more information regarding Initial setup of the CMM, see *Implementing Systems Management of IBM PureFlex System, SG24-8060*.

3. After a successful login, CMM opens the default view showing Chassis Map, as shown in Figure 3-6.

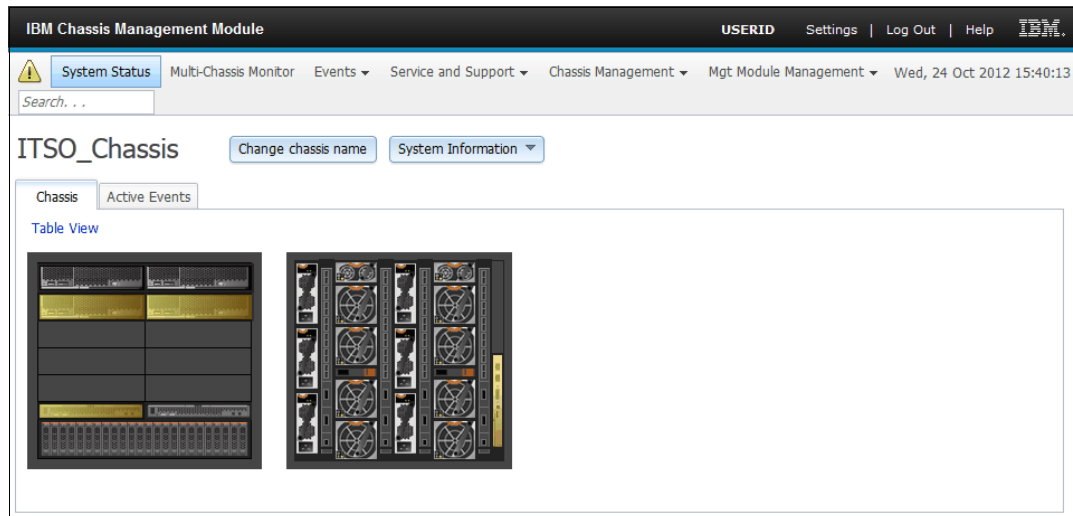


Figure 3-6 IBM Flex System Chassis Management Module Chassis Map (default) view

Disclaimer: In Figure 3-7, and also in some preceding figures, you might find a *yellow* exclamation point on the top left of the CMM interface along with yellow boxes around some components indicating a error condition. The reason for this visual discrepancy is that the various components of the chassis were being tested during the time of writing of this book.

3.2.3 Viewing and configuring IP addresses of chassis components

Due to the increasing demand for Remote Access to IT equipment over networks, being able to easily view and configure IP addresses of devices in a solution is of a great advantage to IT Administrators and Users. The IBM Flex System IBM Flex System Chassis Management Module gives options to conveniently find the IP information of all chassis components in a single view, and gives the ability to easily configure the networking options for a device.

- ▶ In IBM Flex System Chassis Management Module web-interface, a dedicated option is given in the top-level menu item: **Chassis Management** → **Component IP information** as shown in Figure 3-7.

Chassis Management ▾	Mgt Module Management ▾	<input type="text" value="Search. . ."/>
Chassis	Properties and settings for the overall chassis	
Compute Nodes	Properties and settings for compute nodes in the chassis	
Storage Nodes	Properties and settings for storage nodes in the chassis	
I/O Modules	Properties and settings for I/O Modules in the chassis	
Fans and Cooling	Cooling devices installed in your system	
Power Modules and Management	Power devices, consumption, and allocation	
Component IP Configuration	Single location for you to view and configure the various IP address setting of chassis components	
Chassis Internal Network	Provides internal connectivity between compute node ports and the internal CMM management port	
Hardware Topology	Hierarchical view of components in your chassis	
Reports	Generate Reports of hardware information	

Figure 3-7 Component IP configuration

- ▶ A window opens, giving a single location to view the IP addresses of various chassis Components. If you want to view the IP address of a certain component instantly, you can just scroll the mouse pointer over the View option. This action opens a pop-up window that shows the IPv4 and IPv6 addresses of that device, as shown in Figure 3-8.

Component IP Configuration
Configure IPv4 and IPv6 address information for the components below.

I/O Modules

Bay	Device Name	IPv4 Enabled	IP Address
1	Ethernet ScSE	Yes	View
2	Ethernet ScSE	Yes	View
3	IBM FC3171 8Gb SAN Switch	Yes	View
4	IBM FC3171 8Gb SAN Switch	Yes	View

Compute Nodes

Bay	Device Name	IPv4 Enabled	IP Address
11	node11	Yes	View
12	node12	Yes	View
13	node13	Yes	View
14	node14	Yes	View

Storage Nodes

Bay	Device Name	IPv4 Enabled	IP Address
1-4:1	Storage ITE Mgmt	Yes	View
1-4:2	node02	Yes	View

IPv4 Addresses
9.37.117.149
9.37.117.130
9.37.117.150

IPv6 Addresses
fd8c:215d:178e:c0de:5ef3:fcff:fe6f:5cc
fe80::5ef3:fcff:fe6f:5cc

Close

Figure 3-8 Component IP Configuration view

- ▶ To configure the IP address of a certain chassis component, click the component (*under Device Name*), as shown in Figure 3-9.

IP Address Configuration Storage ITE Mgmt

General Setting | **IPv4** | IPv6

Current IP Configuration

Network Interface: eth0

Configuration Method: Use Static IP Address

IP Address: 9.37.117.149

Subnet Mask: 255.255.255.0

Gateway Address: 9.37.117.1

Change IP Configuration

Configuration Method: Use Static IP Address

New Static Address Information

IP Address: 9.37.117.149

Subnet Mask: 255.255.255.0

Gateway Address: 9.37.117.1

Apply

Close

Figure 3-9 IP Address Configuration

You can also view and configure IP addresses using IBM Flex System Manager as well as the node-specific management controllers.

There are numerous functions and features offered by the IBM Flex System Chassis Management Module that are beyond the scope of this publication. However, we cover the I/O module access features in the next section, due to its relevance with storage management.

For detailed information regarding IBM Flex System Chassis Management Module configuration capabilities, visit the following website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.cmm.doc/cmm_ui_config_the_cmm.html

Or, see *Implementing Systems Management of IBM PureFlex System*, SG24-8060.

3.2.4 Accessing I/O modules using CMM

The IBM Flex System Chassis Management Module performs end-to-end management of all components in the solution. It allows many functions and features that can be performed on I/O Modules, as well as, it gives access to platform-based management interfaces of the modules for in-depth management.

To access the properties and configuration of I/O Modules through IBM Flex System Manager, click **Chassis Management** → **I/O Modules**, as shown in Figure 3-10.

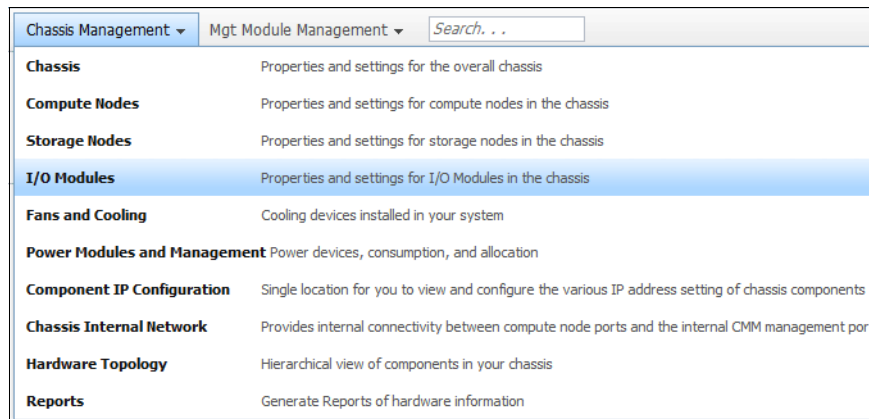


Figure 3-10 FSM - Accessing I/O Modules

The next view shows the available I/O modules in IBM Flex System. It also gives options to Power and Restart the modules as well as the Actions that can be performed on them, as shown in Figure 3-11.

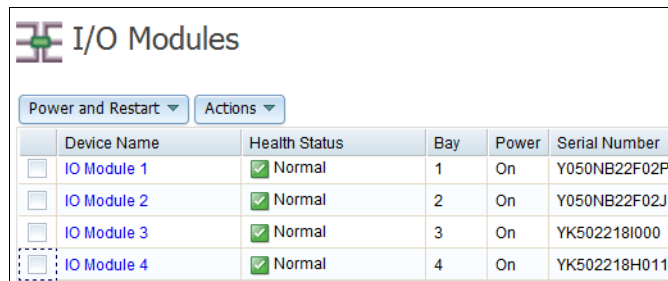


Figure 3-11 Viewing I/O Modules in FSM

Clicking an I/O Module (*Device Name*) opens up a window showing its properties (Figure 3-12).

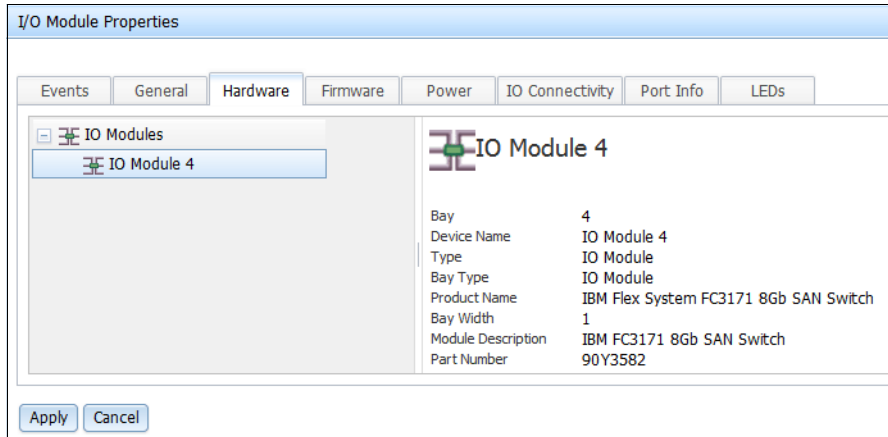


Figure 3-12 Properties of I/O Module device

The Properties window has multiple tabs, showing detailed information regarding the module. Click the **IO Connectivity** tab to view the connectivity details between the I/O module and the compute nodes in the IBM Flex System, as shown in Figure 3-13.

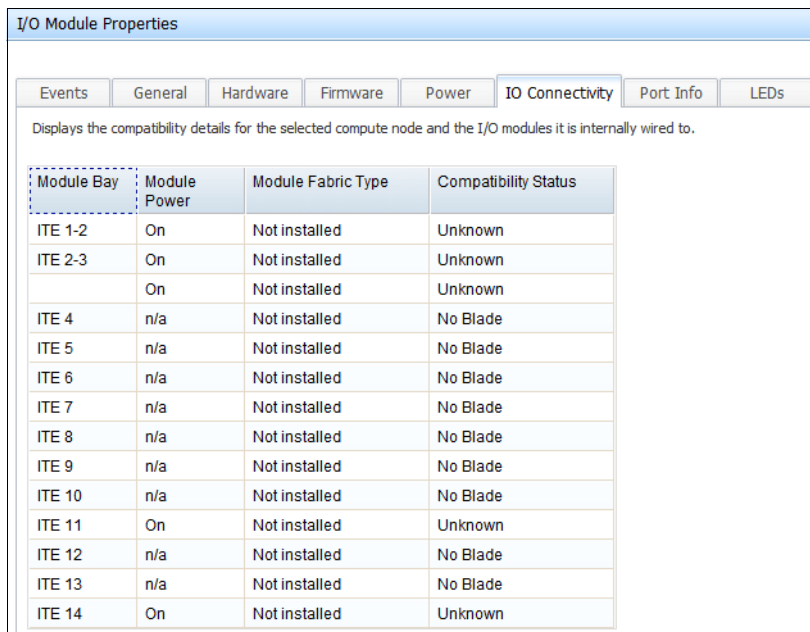


Figure 3-13 IO Connectivity Status of an I/O Module

To open the Module-specific Management interface, click the **I/O Module** wanted, then the **Actions** drop-down button, as shown in Figure 3-14.

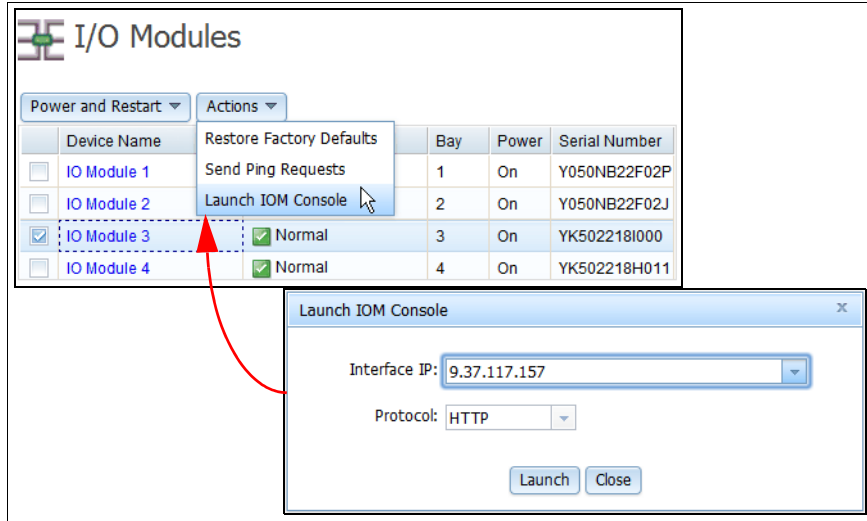


Figure 3-14 Launching IOM Console by CMM

The drop-down menu gives the option to select HTTP or HTTPS protocol. In case of most I/O Modules, HTTPS might need to be selected. It will result in execution of the Switch Modules platform-level management interface.

3.2.5 Managing storage using IBM Flex System Chassis Management Module

After completing initial setup, a subsequent login into CMM opens the *System Status* page. The top of the page identifies the current user and contains a *Settings* menu and *Log Out* option, as shown in Figure 3-15.

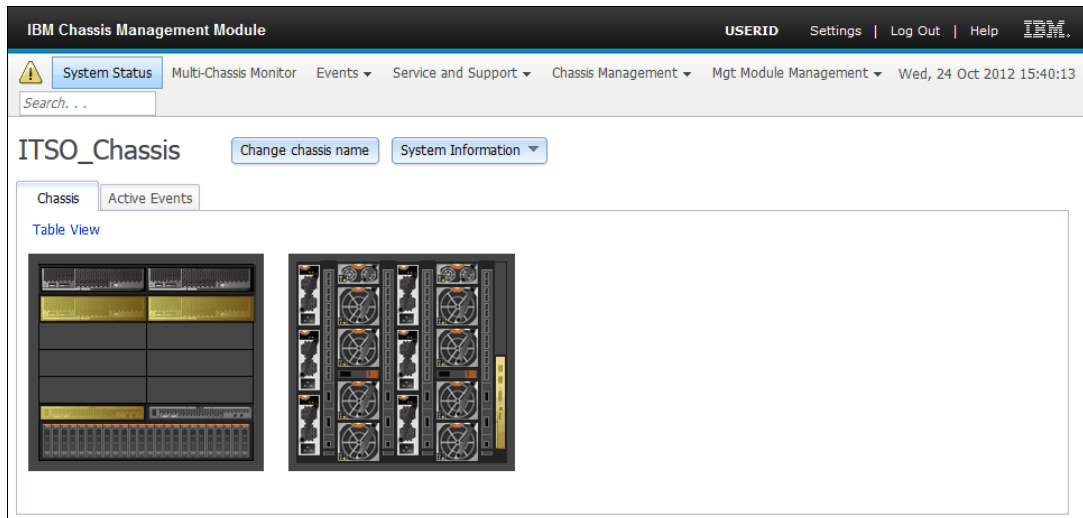


Figure 3-15 IBM Chassis Management Module Chassis Map

The *System Status* shows the view of Chassis, front and back. It also shows health status of the individual physical components.

Tip: Except for the top menu items, the remaining content of the *System Status* page and other web interface pages will vary according to the type of IBM Flex System Enterprise Chassis that you are using and the firmware versions and options that are installed.

Scrolling the mouse on the selected component opens a pop-up windows that lists its name, location, and FRU number, as shown in Figure 3-16.

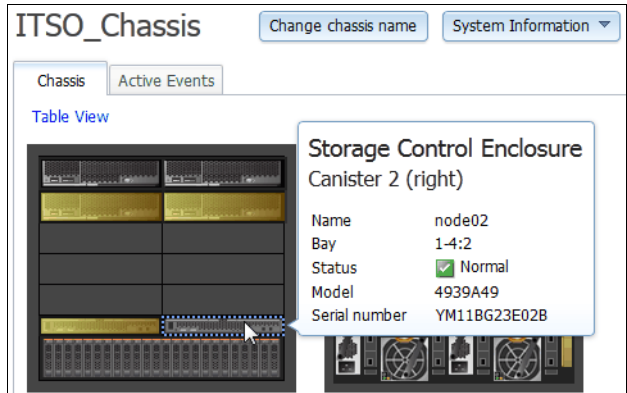


Figure 3-16 CMM chassis view - component pop-up

When you click a component, details about it emerge in the lower part of interface. Depending on the component, one or more tabs show information regarding hardware, firmware, events, and other information. An example is shown in Figure 3-17.

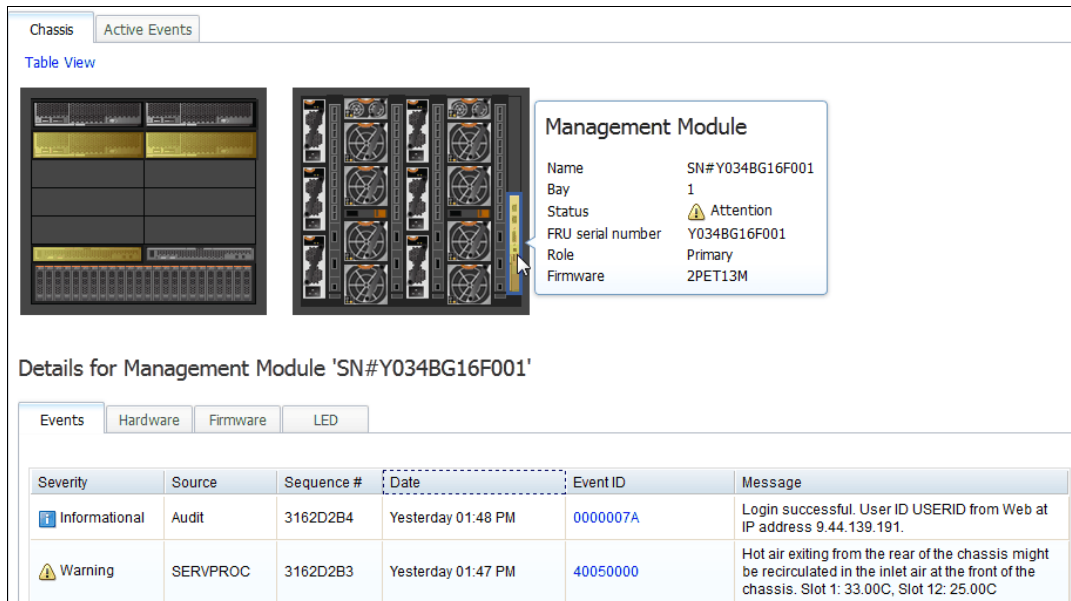


Figure 3-17 CMM - Details of the selected Chassis component

The *System Information* (blue) button, which is located above the graphical representation of the Chassis, gives a quick view of the selected component's information.

The quick view includes part information, serial number, health and temperature, and unique identity (depending on the component), as shown in Figure 3-18.

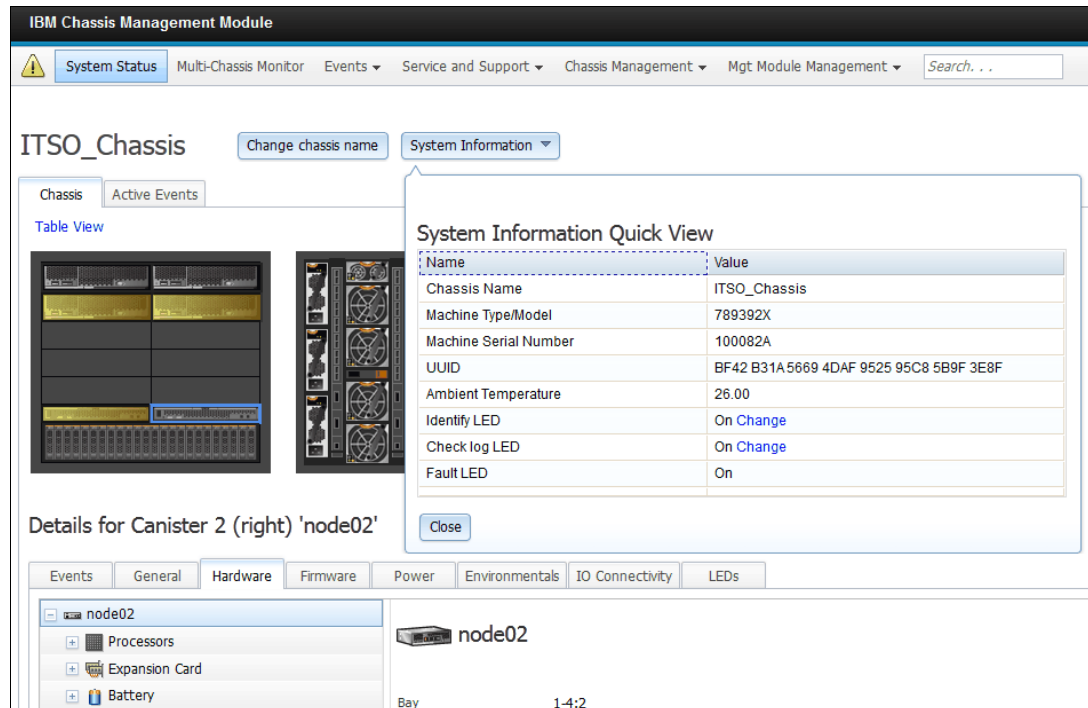


Figure 3-18 CMM - System Information Quick View

Selecting a component also opens up an Actions frame, which lists the actions that can be performed on that particular IBM Flex component, as shown in Figure 3-19.

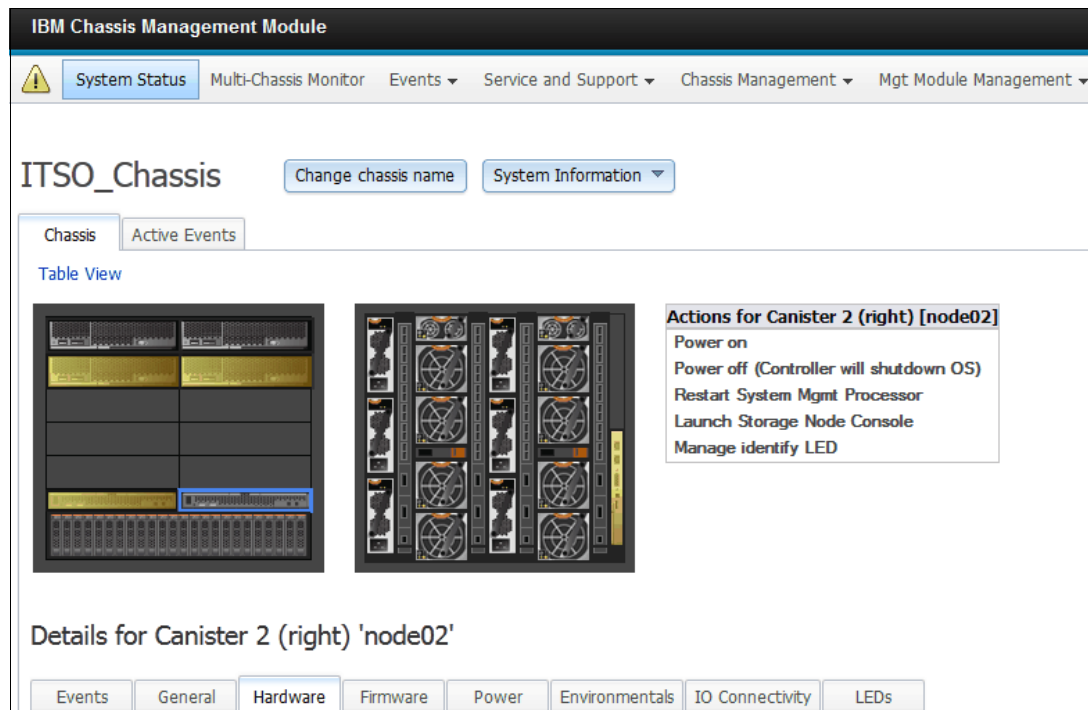


Figure 3-19 CMM - Actions for Chassis Component

In Figure 3-19 on page 98, an IBM Flex System V7000 Storage Node canister is seen as selected (blue rectangle around component). On the right-side of the Chassis rear view, there is a list of Actions for the particular Canister, which includes options to power on/off, restart the service processor, launch the detailed IBM Flex System V7000 Management application, as well as manage the identity LEDs on the component.

The area of GUI under the graphical representation of the chassis shows details regarding the selected component. This includes very detailed information laid out in convenient tabular format. The example in Figure 3-20 shows firmware information including its type, build information, release date, revision, and role.

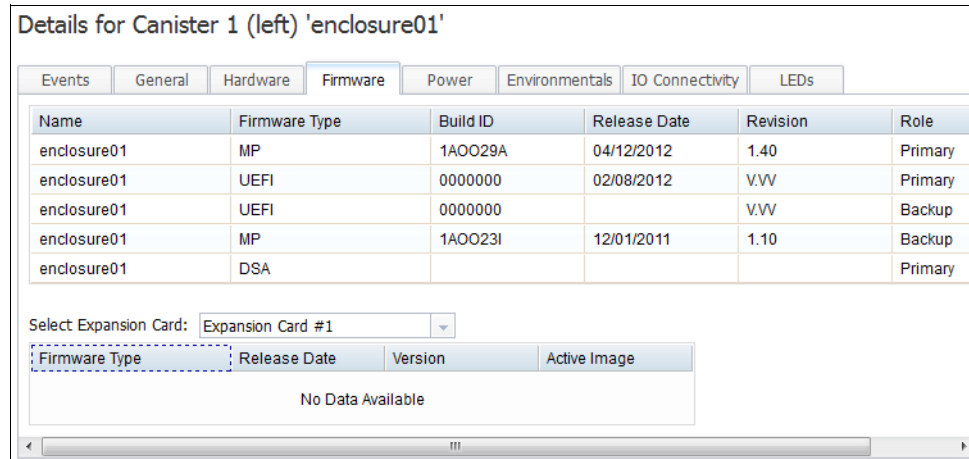


Figure 3-20 Details regarding selected chassis component - Firmware tab

With the advent of complex integrated IT solutions, sophisticated hardware with varying power and voltage requirements get consolidated into a typical enterprise rack. To ensure proper functioning of all the hardware, IT administrators have to ensure that all equipment stays below the temperature, power, and voltage thresholds at all times. The IBM Flex System Chassis Management Module web interface provides multiple views to facilitate monitoring and notification of environmental and temporal indices.

Figure 3-21 shows the Environmental details of various components within selected chassis hardware (storage node canister in this case). It shows the current values as well as the thresholds set for these components.

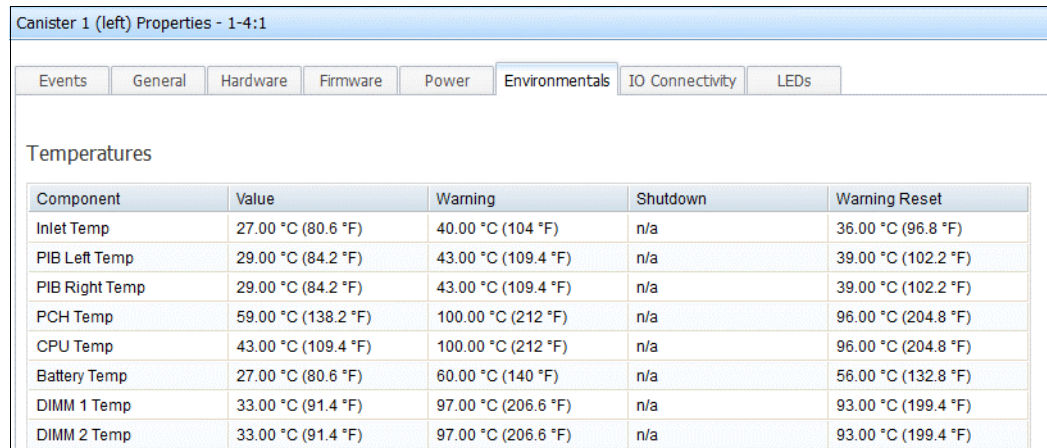


Figure 3-21 Canister properties - Environmental

The Environmental view also reports voltage value and thresholds, as shown in Figure 3-22.

Voltages		
Source	Value	Critical
Planar 3.3V	+3.36	(+2.96,+3.62)
Planar 5V	+5.12	(+4.48,+5.49)
Planar 12V	+12.43	(+10.78,+13.20)
Planar VBAT	+2.95	(+2.70,+3.18)
VRM1	+115.75	n/a

Figure 3-22 Canister properties - Voltages

An important view for the storage canister is the I/O Connectivity tab, which shows the fabric information related to the I/O modules in the chassis, as shown in Figure 3-23.

Events	General	Hardware	Firmware	Power	Environmentals	IO Connectivity	LEDs
Displays the compatibility details for the selected compute node and the I/O modules it is internally wired to.							
Module Bay	Module Power	Module Fabric Type	Fabric On Compute Node	Compatibility Status			
IOM 1	On	Ethernet ScSE	Unknown	OK			
IOM 2	On	Ethernet ScSE	Unknown	OK			
IOM 3	On	Fibre ScSE	Unknown	OK			
IOM 4	On	Fibre ScSE	Unknown	OK			
<input type="button" value="Apply"/> <input type="button" value="Cancel"/>							

Figure 3-23 IO Connectivity

Figure 3-24 shows the DC Power consumption for the component.

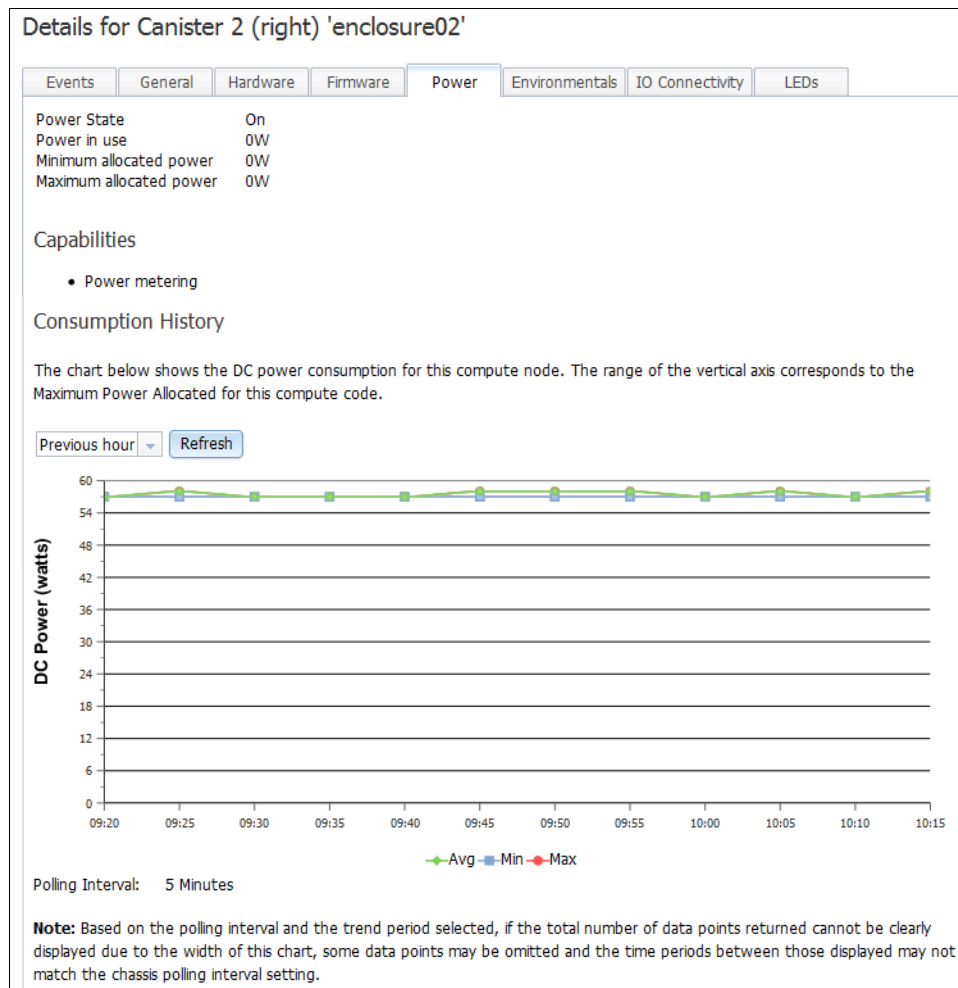


Figure 3-24 Power consumption of selected chassis component

The IBM Flex System Chassis Management Module web interface has a menu structure at the top of each page that gives easy access to most management functions. One of the most frequently used menu items, Chassis Management, shows properties, settings, and various views and reports regarding the nodes, modules, and other components in the chassis.

Figure 3-25 shows the option that gives the hierarchical view of chassis components.

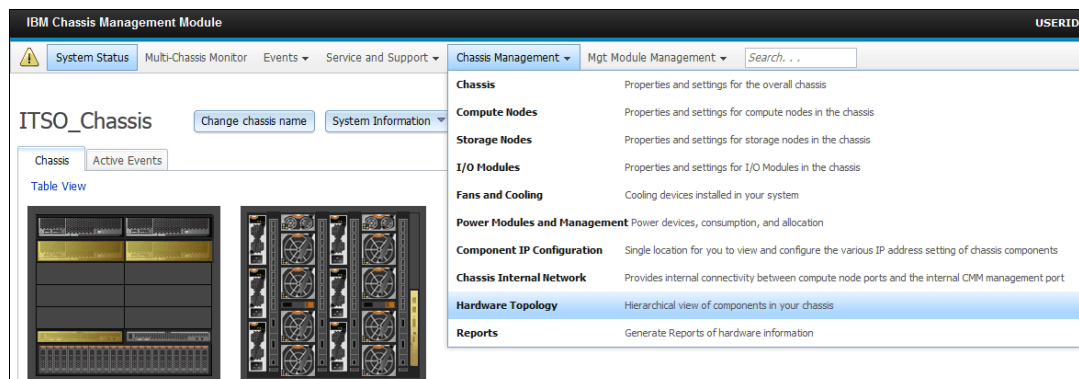


Figure 3-25 Hardware Topology view from Chassis Management menu

The Chassis Hardware Topology view provides very detailed information regarding all the hardware in the chassis. It goes deep to the point where you can select individual drives, batteries, and even expansion cards within a storage or compute node, as shown in Figure 3-26.

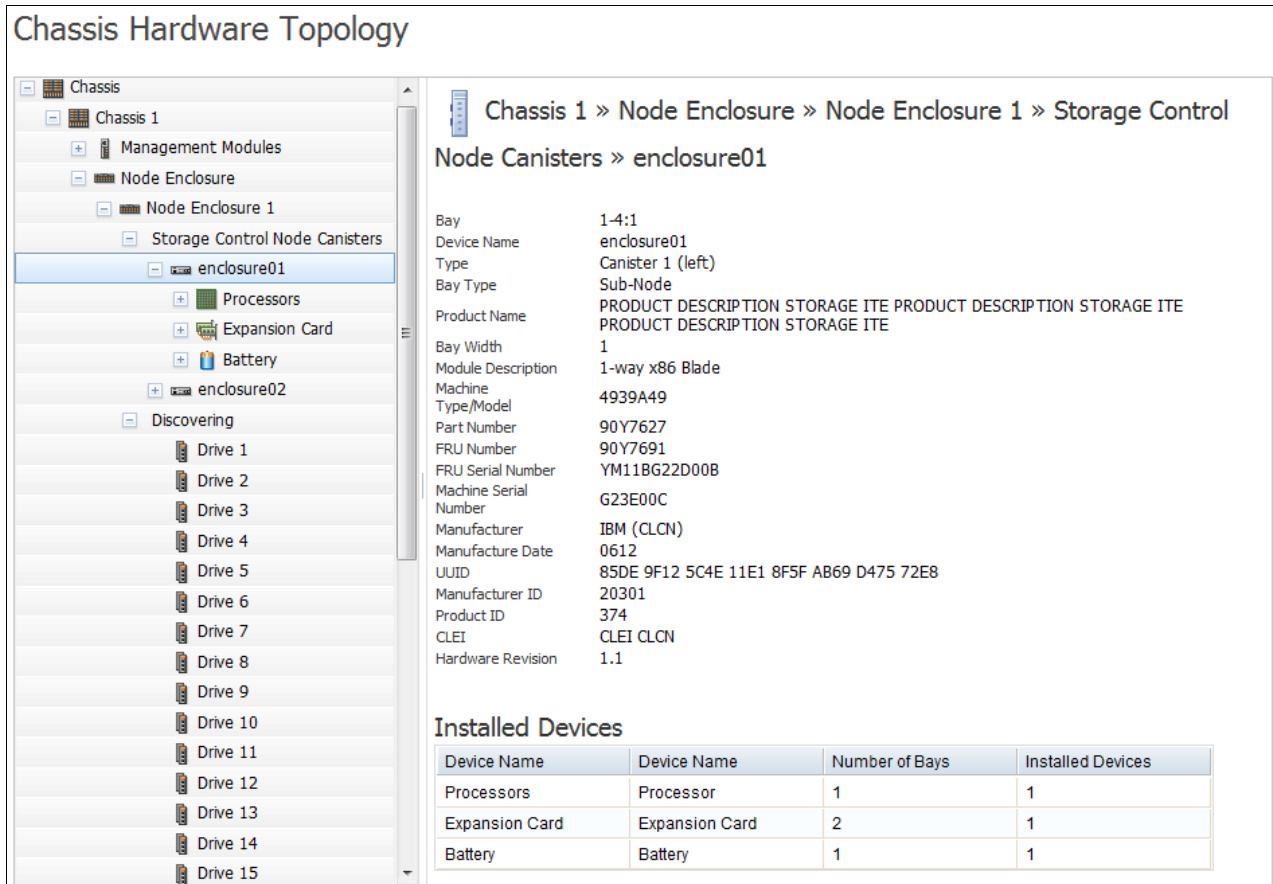


Figure 3-26 Chassis Hardware topology

- From the top menu, select **Chassis Management** → **Storage Nodes** to view details regarding storage nodes, as shown in Figure 3-27.

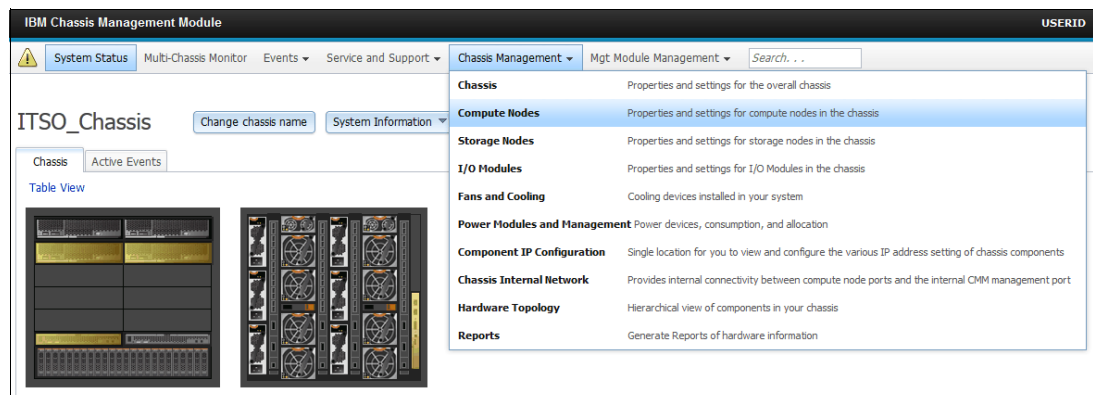
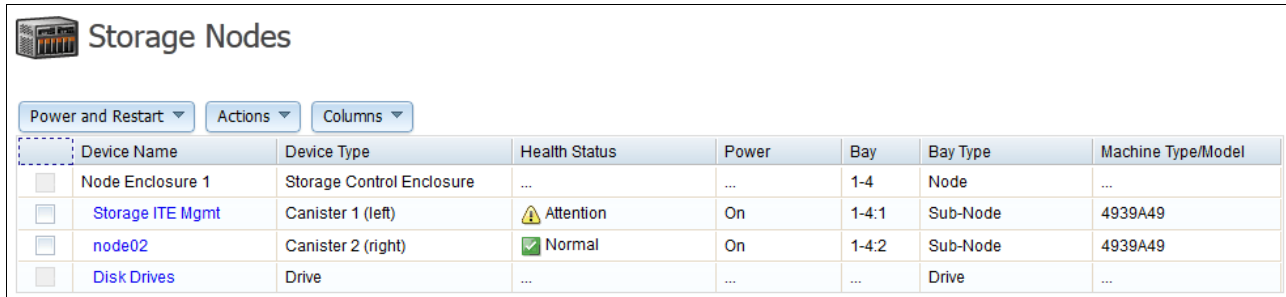


Figure 3-27 Chassis Management - Storage Nodes

The Storage Nodes view shows all the storage enclosures in the chassis, and gives options to perform various actions on these, including power and restart options, as shown in Figure 3-28.

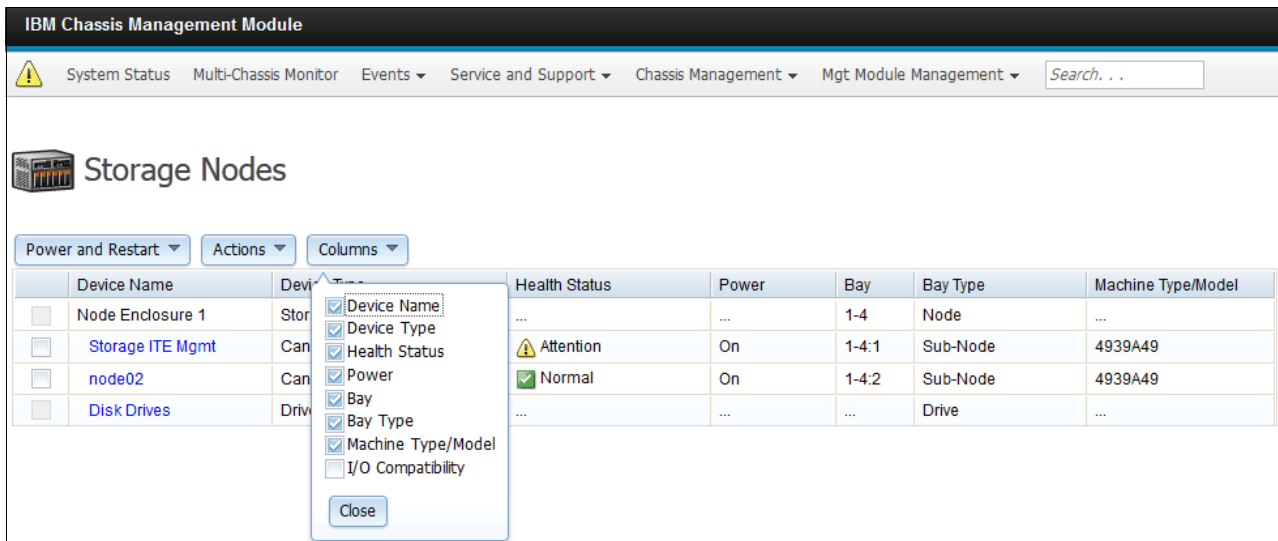


The screenshot shows the 'Storage Nodes' view with a table containing the following data:

Device Name	Device Type	Health Status	Power	Bay	Bay Type	Machine Type/Model
Node Enclosure 1	Storage Control Enclosure	1-4	Node	...
Storage ITE Mgmt	Canister 1 (left)	Attention	On	1-4:1	Sub-Node	4939A49
node02	Canister 2 (right)	Normal	On	1-4:2	Sub-Node	4939A49
Disk Drives	Drive	Drive	...

Figure 3-28 CMM - Storage Nodes - view of all enclosures

In the Storage Node view, you can select the column headers depending on how much detail is wanted, as shown in Figure 3-29.



The screenshot shows the 'Storage Nodes' view with the 'Columns' dropdown menu open. The menu options are:

- Device Name
- Device Type
- Health Status
- Power
- Bay
- Bay Type
- Machine Type/Model
- I/O Compatibility

The table data is the same as in Figure 3-28.

Figure 3-29 CMM - Storage Nodes view - columns

The Actions menu gives options to manage identify LED of the selected enclosure, as shown in Figure 3-30.

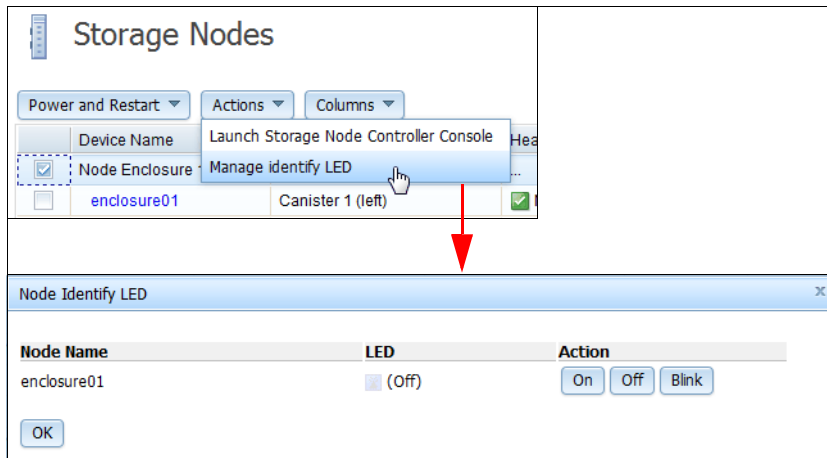


Figure 3-30 Manage Identify LED

The IBM Flex System V7000 Storage Node User Interface is used for detailed-level management and configuration of the IBM Flex System V7000 Storage Node.

Click **Actions** → **Launch Storage Node Controller Console**, which will take you to the IBM Flex System V7000 Storage User Interface, as shown in Figure 3-31.

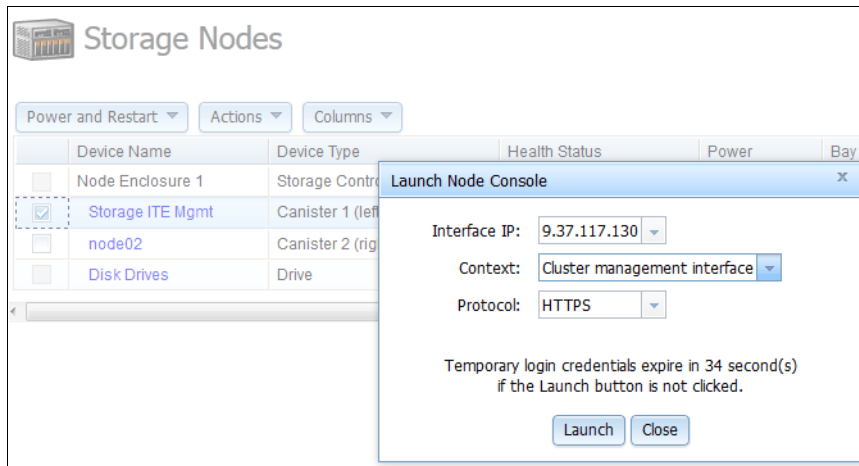


Figure 3-31 CMM - Launch Storage Node Controller Console

The next window is the IBM Flex System V7000 Storage Management application's login panel, as shown in Figure 3-32.



Figure 3-32 IBM Flex System V7000 Storage Node Login window

After logging in, you can use the IBM Flex System V7000 Storage User Interface directly to perform the management tasks, as shown in Figure 3-33.

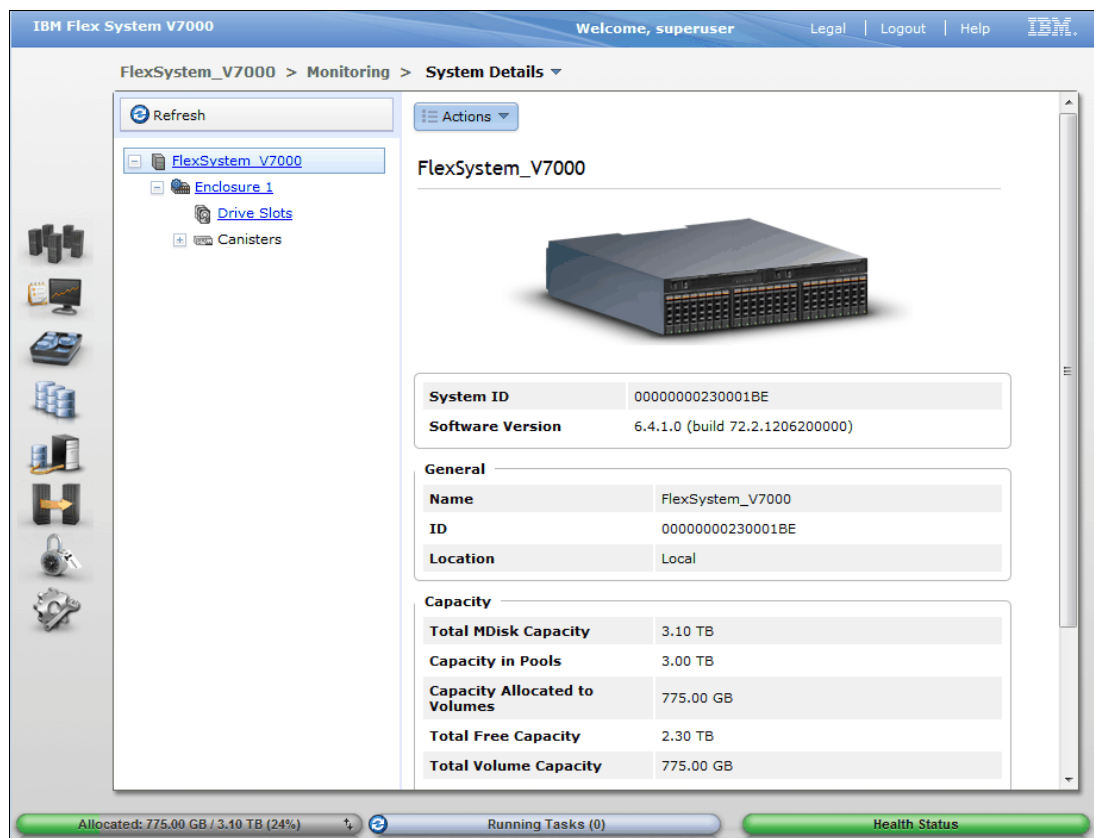


Figure 3-33 IBM Flex System V7000 System Details

For further information regarding IBM Flex System V7000 Storage User Interface, see Chapter 5, “IBM Flex System V7000 Storage Node GUI interface” on page 175.

3.2.6 Data collection using CMM

The IBM Flex System Chassis Management Module offers various powerful options to pull service information from the components in chassis. The *Service and Support* menu allows the user to view a list of serviceable problems and their current status, enable and configure electronic Call Home capability, review detailed status about CMM's health and connectivity, and generate as well as retrieve service data files from the chassis components.

To acquire service data from the Storage Node, you can perform the following steps:

- ▶ From the top menu, select **Service and Support** → **Download Service Data**, as shown in Figure 3-34.

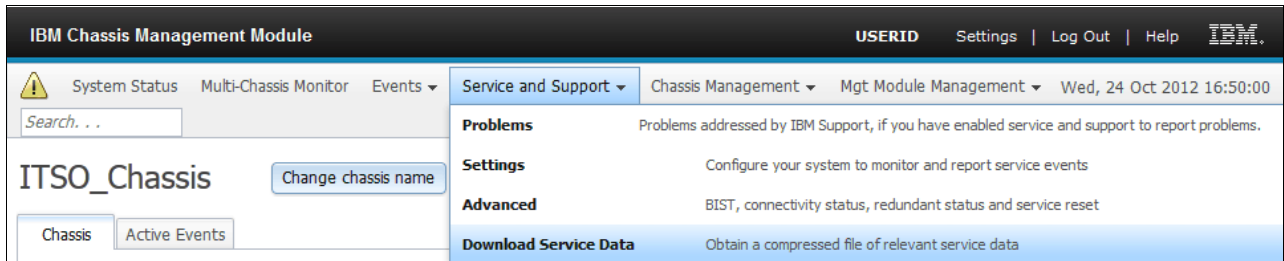


Figure 3-34 CMM - Service and Support - Download Service Data

- ▶ The *Download Service Data* page gives four options. We pick the option, **Storage Node Service Data**, as shown in Figure 3-35.




Figure 3-35 CMM - Download Service Data from Storage Node

The next window gives three options to dump the Storage Node Service Data. A description of all three options is also provided, so the user can easily choose the option wanted.

- ▶ In our case, we click the drop-down menu **Dump Type:** and then select **Service Data**. Choose the option **Initiate and Collect**, which allows the storage node to initiate the generation of a new dump file on the compute node and overwrite any existing file. The CMM will then collect the dump file by transferring it from the node to the CMM file system, as shown in Figure 3-36.

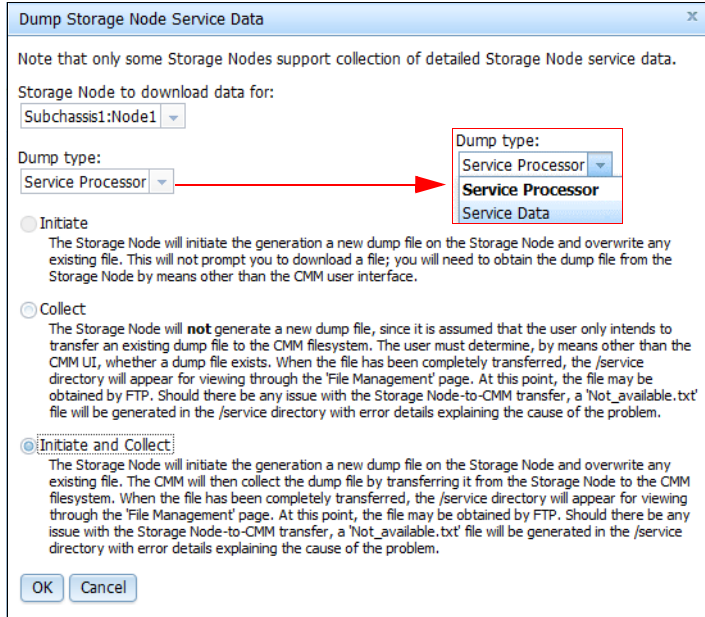


Figure 3-36 Initiate and Collect

- ▶ The system will prompt you for verification, indicate that operation is in progress, and then its completion, as shown in Figure 3-37.

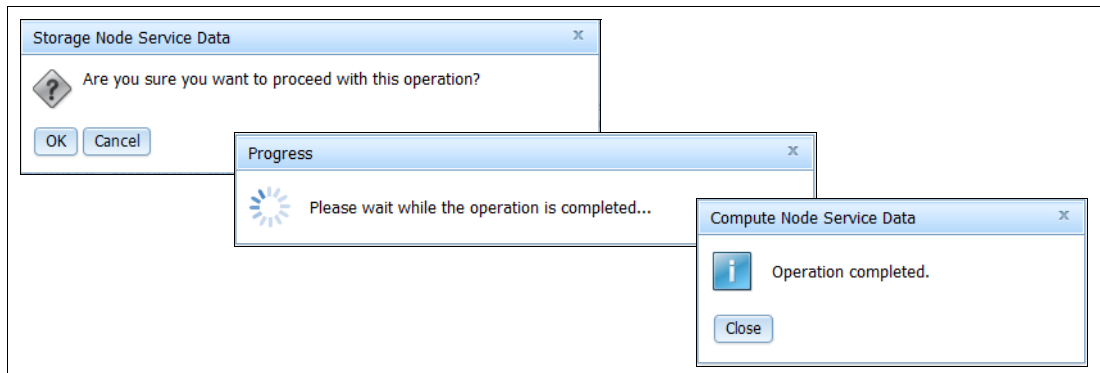


Figure 3-37 Status dialog boxes

To view or delete files in the CMM local storage file system, select **Mgt. Module Management** → **File Management**, as shown in Figure 3-38.

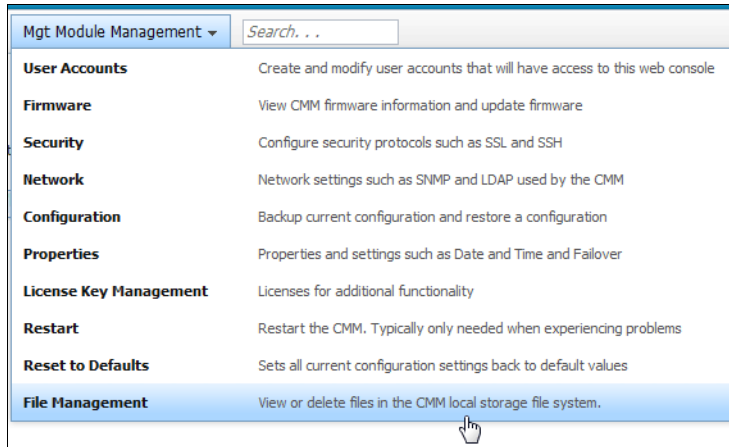


Figure 3-38 File Management

The File Management page shows the status of CMM's local storage file system and the CMM File Management menu, as shown in Figure 3-39.

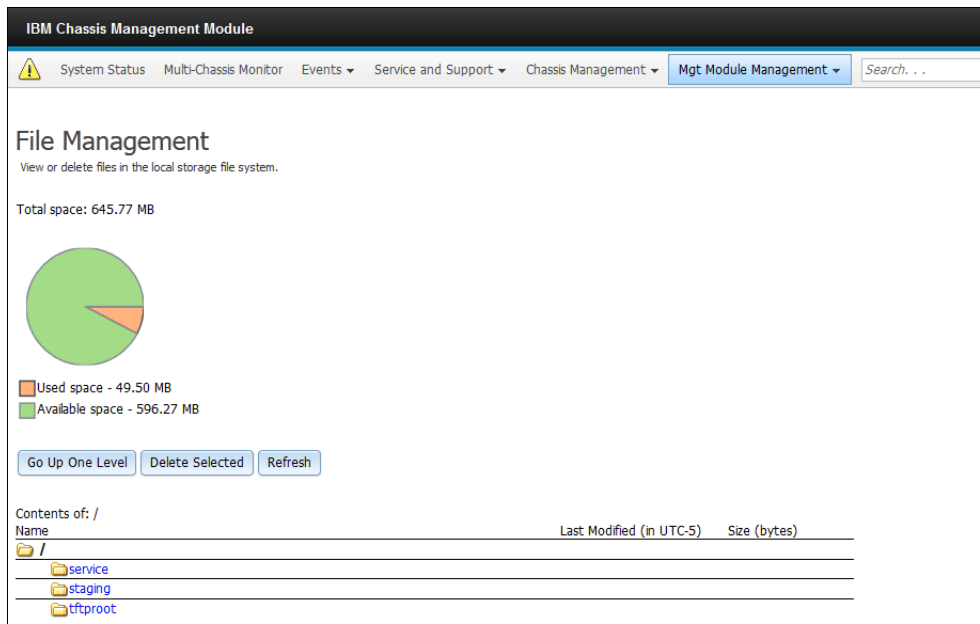


Figure 3-39 Management of local CMM file system

The /service/ folder contains the file that contains the data dump, as shown in Figure 3-40.

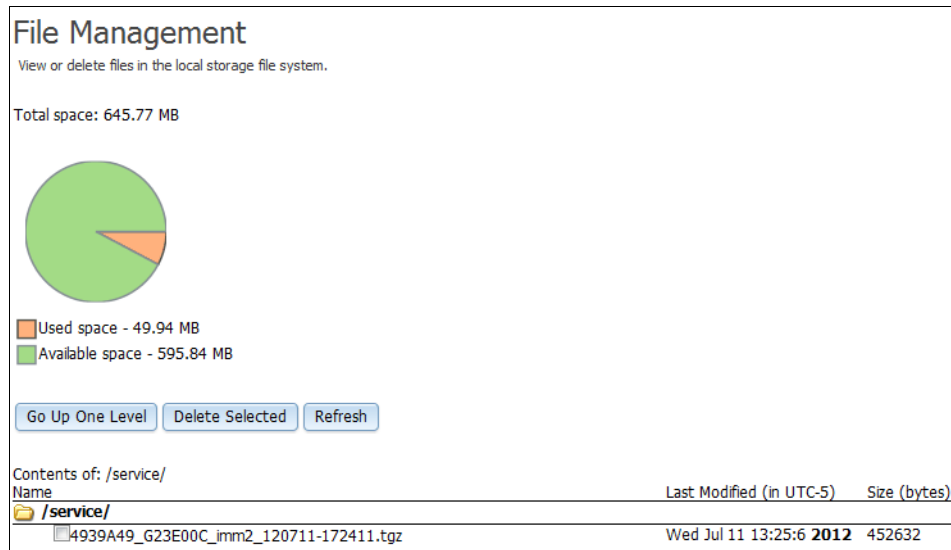


Figure 3-40 Viewing service data in /service/ folder

The file containing a dump of service data is also copied in /tftpboot/service/ for ease of transfer, as shown in Figure 3-41.



Figure 3-41 Viewing service data in /tftp/service/ path

For more details regarding the IBM IBM Flex System Chassis Management Module, visit the following website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.cmm.doc/cmm_product_page.html

3.3 Flex System Manager (FSM)

The IBM Flex System Manager Node (FSM) is an advanced chassis management appliance that provides the ability to manage up to four chassis. Also referred to as the Management Node, FSM is based on an Intel-based x86 compute node that comes with preloaded management software. FSM software contains a fully operational virtual infrastructure management stack for managing multiple chassis from a single focus-point.

IBM Flex System Manager provides an instant and resource-oriented integrated view of chassis and chassis resources, including compute and storage nodes as well as I/O modules. It simplifies management by facilitating faster implementation, reconfiguration and provisioning of server, storage and network resources. It acts as a single-point of support across all assets and resources in IBM PureFlex Systems for the complete range of management tasks, including service requests, incident handling and problem resolution, as well as handling specialized management tasks in connection with security, performance, resource allocation, and application version control, and so on.

Figure 3-42 shows the front view of an IBM Flex System Manager Node.



Figure 3-42 IBM Flex System Manager Node

Tip: The IBM Flex System Chassis Management Module provides control and information on the storage hardware, while the IBM Flex System Manager additionally provides configuration functionality.

3.3.1 Overview of IBM Flex System Manager (FSM)

The comprehensive management capabilities of IBM Flex System Manager (FSM) allows for seamless and simplified storage management, virtualization, and networking.

Consolidating management and administration tasks into a single, consistent interface while automating many tasks across various compute and storage nodes simplifies the job of administrators to manage and configure servers, storage, network, and other resources.

Whole rack management of energy, power, cooling, and so on, reduces cost and cuts downtime, while the systems operate at higher performance levels and are more consistent in operation. The single interface eliminates the need of having different administrative support and skills for each platform, thus reducing training costs.

The IBM Flex System supports a secure policy by default to ensure a secure chassis infrastructure, and a Legacy policy that offers flexibility in chassis security. All components run the same security policy. A centralized security policy makes Enterprise Chassis easy to configure. The system password is required by the management software to perform system-management functions, such as enabling access to the local user registry. The management software system password is specified when the management software is set up initially, during the Management Server Setup wizard. For more information regarding FSM initial setup, see *Implementing Systems Management of IBM PureFlex System*, SG24-8060.

IBM Flex System Manager will auto-discover new and existing components, as well as automatically initiate and perform set-up tasks. As part of this discovery process, IBM Flex System Manager prepares and maintains complete, detailed front, back and Chassis Map views of each chassis and all physical compute nodes and virtual servers that exist in the rack.

Management and control of the physical and virtualized infrastructure includes initiating the detailed tasks required for full life-cycle management of virtualized devices, including their creation, provisioning, configuration, operation and de-provisioning. It can build and maintain a library of pre-defined virtual devices. This allows fast automated, on-demand adjustments in server, storage and network capacities in response to changing business and operational needs.

The built-in network provides a reliable and secure communications path for management and inter-rack data sharing. Sophisticated network resource management moves the focus from managing the network to managing the workload. The complete infrastructure is managed to assure maximum workload performance. This allows for operational environment-aware automated smart placement and provisioning of the workload itself. The emphasis is on optimized, automated, integrated network control and management.

Here are various other significant characteristics of an IBM Flex System Manager:

- ▶ All basic and advanced functions come preloaded as an appliance.
- ▶ A real-time, multi chassis view of hardware components is presented with overlays for additional information.
- ▶ Easy-to-use multi-chassis management capabilities are provided:
 - Quick Start Wizards with automated discovery
 - Advanced remote presence console across multiple chassis
 - Centralized IBM Feature on Demand (FoD) license key management
- ▶ Integrated x86 and Power nodes, Storage, and Network management.
- ▶ Virtualization Management including resource pools.
- ▶ Robust Security (Centralized User Management, Security Policy, Certificates).
- ▶ Integrated LDAP and NTP servers for private management network.
- ▶ Active Energy Management.
- ▶ Active/Passive Failover to redundant FSM Appliance.
- ▶ Additional advanced management capabilities available through add-ons.
- ▶ Upward integration into Systems Director, Tivoli, other Third Party Enterprise Managers.
- ▶ Configuration Manager plug-ins (besides IBM Flex System Chassis Management Module):
 - Configuration of Chassis and Internal Switch Components
 - Configuration of Compute Nodes

- ▶ Remote Presence:
 - Login and take over Compute node from Administrators console
 - Mount DVD drive remotely
 - Perform Bare Metal Install remotely
- ▶ Service and Support Management with Call-Home:
 - Automatically detect and log HW problems
 - Notify IBM directly of HW Outages
- ▶ Required update of the default passwords after the initial setup.
- ▶ Only secure communication protocols, such as SSH and SSL. Unencrypted protocols, such as HTTP, Telnet, and SNMPv1, are disabled.
- ▶ Certificates to establish secure, trusted connections for applications that run on the management processors.
- ▶ Strong password policies with automatic validation and verification checks.

For further details regarding IBM Flex System Manager features and capabilities, see *Implementing Systems Management of IBM PureFlex System, SG24-8060*.

3.3.2 IBM Flex System Manager storage management features

IBM PureFlex offers interoperability with a wide range of storage subsystems.

IBM Flex System Manager overview

The system infrastructure has the following salient features:

- ▶ IBM Flex System V7000 Storage Node: Integrated virtualized IBM Flex System Storage
- ▶ IBM Flex System Storage Virtualization: Used to virtualize external storage for greater data center efficiency and utilization (such as IBM Storwize V7000, IBM Flex System V7000 Storage Node, and IBM SAN Volume Controller)
- ▶ IBM Flex System Storage Interoperability: A broad set of IBM storage is supported with IBM Flex System. Interoperability with 3rd party storage is also offered via virtualization
- ▶ IBM Flex System FSM Storage Control, which offers the following:
 - Discovery and Inventory
 - Monitoring and Alerts
 - Configuration
 - Provisioning
- ▶ IBM Flex System FSM Advanced, which offers the following features and functions:
 - Integrated virtualization management across server, storage, network
 - Image repository and management Storage provisioning for image creation, deployment, and cloning
 - System Pools
 - Integrated management of storage in life cycle of defining and managing system pools
 - Virtual Image Cloning
 - Integrated storage provisioning and virtual image placement for new virtual machines

The interoperability is illustrated in Figure 3-43.

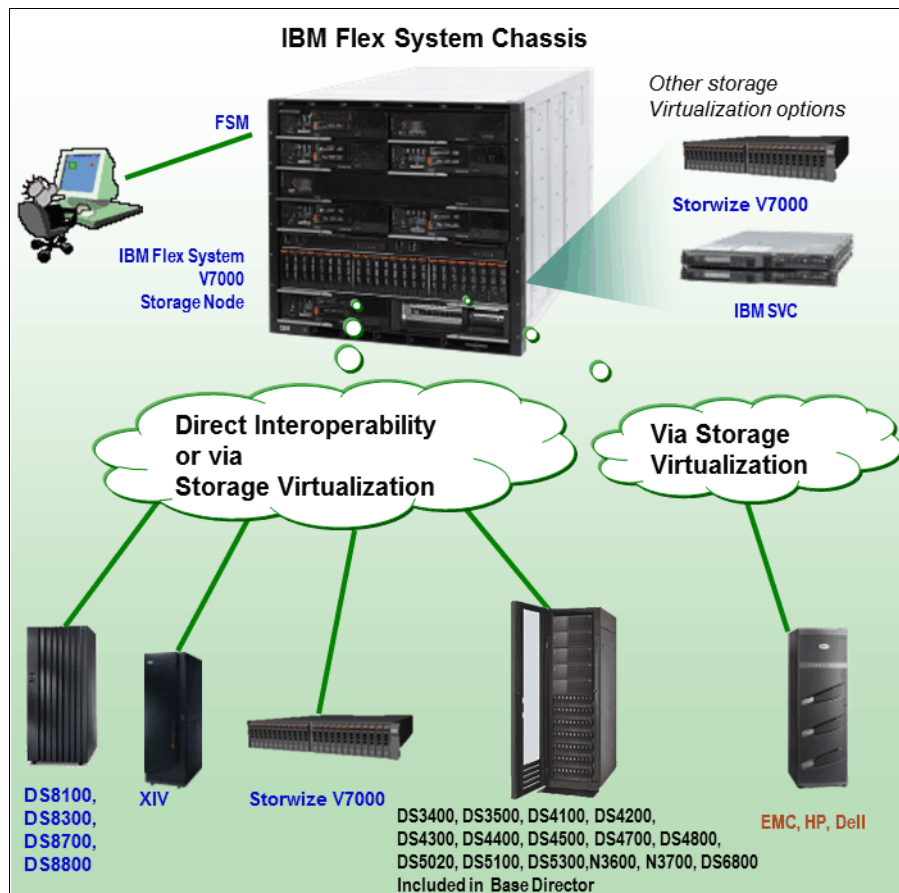


Figure 3-43 Storage Interoperability

IBM Flex System Manager is the common management interface used in an IBM Flex System for multi-chassis management. It offers features for simplified administration of storage along with server and network management. The integrated storage management with Flex System Manager offers simplified storage virtualization capabilities.

The following storage features are included in IBM Flex System Manager:

- ▶ Storage device discovery
- ▶ Showing relationships between storage and server resources
- ▶ Physical and logical topologies
- ▶ Configuration of physical and logical storage
- ▶ Viewing controller and volume statuses
- ▶ Setting notification alerts
- ▶ Integration with VM Control provides:
 - Storage provisioning for virtual server image creation, deployment, and cloning
 - Storage system pool lifecycle management, provisioning, and cloning actions within the pool

You can virtualize and manage the complete solution infrastructure using the IBM Flex System Manager. Figure 3-44 shows an overview of FSM virtualization.

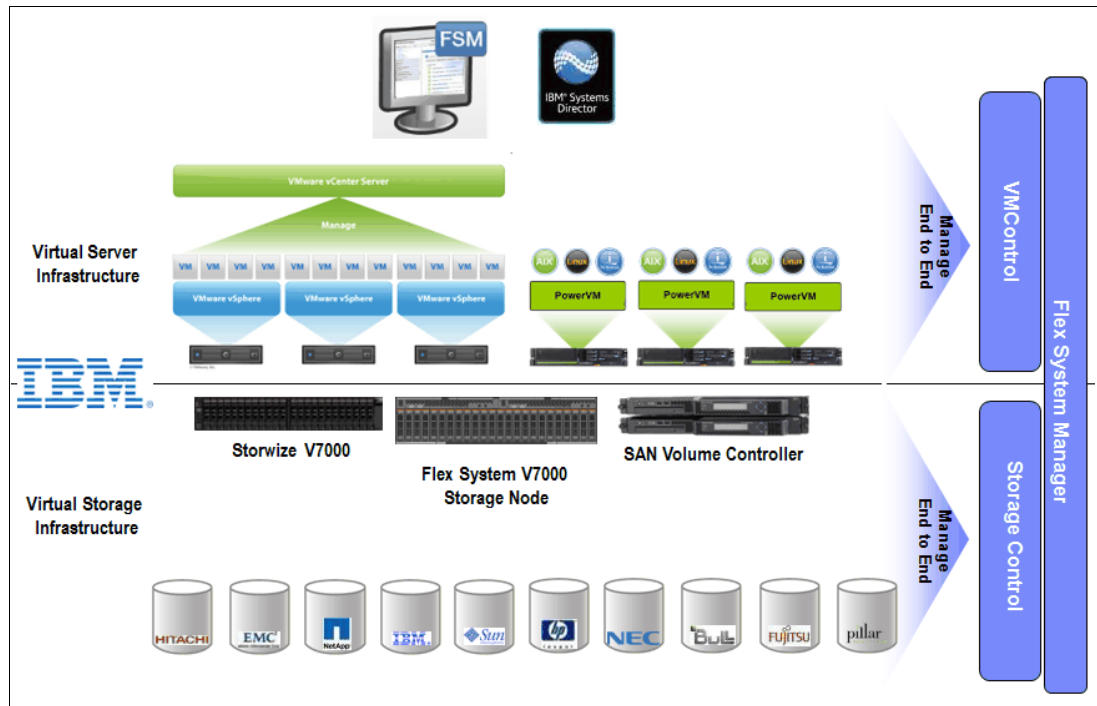


Figure 3-44 Managing diverse IT infrastructure with storage

For further details regarding managing storage using FSM, see 3.3.7, “Managing storage using IBM Flex System Manager” on page 139.

IBM Flex System Manager editions

The IBM Flex System Manager software contains a set of components that are responsible for performing certain management functions. These components must be activated using the available IBM Feature on Demand (FoD) software entitlement licenses, and they are licensed on a per-chassis basis.

Important: The management node comes standard without any entitlement licenses, therefore, you must purchase a license to enable the required IBM Flex System Manager functionality.

IBM Flex System Manager base feature set

The following functionality and features are available at this time in the IBM Flex System Manager base feature set:

- ▶ Support for up to four managed chassis, and up to 5,000 managed elements
- ▶ Auto-discovery of managed elements
- ▶ Overall health status
- ▶ Monitor and availability
- ▶ Hardware management
- ▶ Security management
- ▶ Administration
- ▶ Network management (Network Control).Storage management (Storage Control)
- ▶ Virtual machine life cycle management (VMControl Express)
- ▶ I/O address management (IBM Fabric Manager)

IBM Flex System Manager advanced feature set

The IBM Flex System Manager advanced feature set offers all capabilities of the base feature set plus the following features:

- ▶ Image management (VMControl Standard)
- ▶ Pool management (VMControl Enterprise)

Important: The advanced upgrade license requires the IBM Flex System Manager base license.

IBM Fabric Manager (IFM)

IBM Fabric Manager (IFM) is available as a Feature on Demand (FoD) through the IBM Flex System Manager management software. It is an easy-to-use server provisioning I/O management tool. Unlike performing deployment steps manually, the parameters are configured automatically. IBM Fabric Manager monitors your data center and automatically responds to server issues, which reduces the risk of server failures and downtime.

IFM offers the following enhanced features and functions:

- ▶ Automatic provisioning of I/O addresses.
- ▶ Pre-defined configurations of I/O parameters, such as MAC addresses for Ethernet, World Wide Names and SAN boot targets - for up to 256 chassis or 3584 compute nodes.
- ▶ Provides I/O parameter and VLAN migration to standby compute nodes in case of compute node failure and executes action plans based on triggered events.
- ▶ Monitors the health of compute nodes, and automatically (without user intervention) replaces a failed compute node from a designated pool of spare compute nodes.
- ▶ Using an enhanced graphical user interface (GUI), creates addresses for compute nodes, saves the address profiles, and deploys the addresses to the slots in the same chassis (or in up to 256 different chassis).

IBM Fabric Manager enables you to quickly replace and recover compute nodes in your environment. This is done by assigning Ethernet MAC, Fibre Channel WWN, and SAS WWN addresses so that any compute nodes plugged into those bays take on the assigned addresses, which enables the Ethernet and Fibre Channel infrastructure to be configured once, and before any compute nodes are connected to the chassis.

After receiving a failure alert, IFM attempts to power off the failing compute node, reads the IFM virtualized addresses and boot target parameters, applies these parameters to the next compute node in the standby pool, and power on the standby compute node.

Additionally, you can create profiles for chassis that have not been installed in the environment by simply associating an IP address to the future chassis.

You must configure the CMM to enable IFM access. The IFM configuration is chassis-based and does not transfer with the physical CMM. When a CMM is moved to a new chassis, it clears out its IFM configuration, and the IFM configuration must be reapplied on the new CMM.

Important: You must set the TCP Command Mode Protocol on the CMM Web interface.

There are minimum hardware and software requirements that the system must meet before you can install or use IFM, which are documented in detail at the following website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.iofm.doc/dw1li_creating_a_configuration_file_automatically.html

For information regarding how to install, configure, and use IBM Fabric Manager (IFM), visit this website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.iofm.doc/dw1li_advanced.html

3.3.3 Logging in to the IBM Flex System Manager Node

The following section gives step-by-step instructions regarding how to log in to the FSM:

1. Open a web browser, and type in the IP address of the FSM. Remember that depending on the security style, it might be necessary to use https protocol.
2. At the Login panel, enter your User name and Password, as shown in Figure 3-45.

Figure 3-45 shows the FSM login page.



Figure 3-45 IBM Flex System Manager login panel

- After a successful login, you are taken to the default Home page of IBM Flex System Manager. Figure 3-46 shows the FSM Home page.

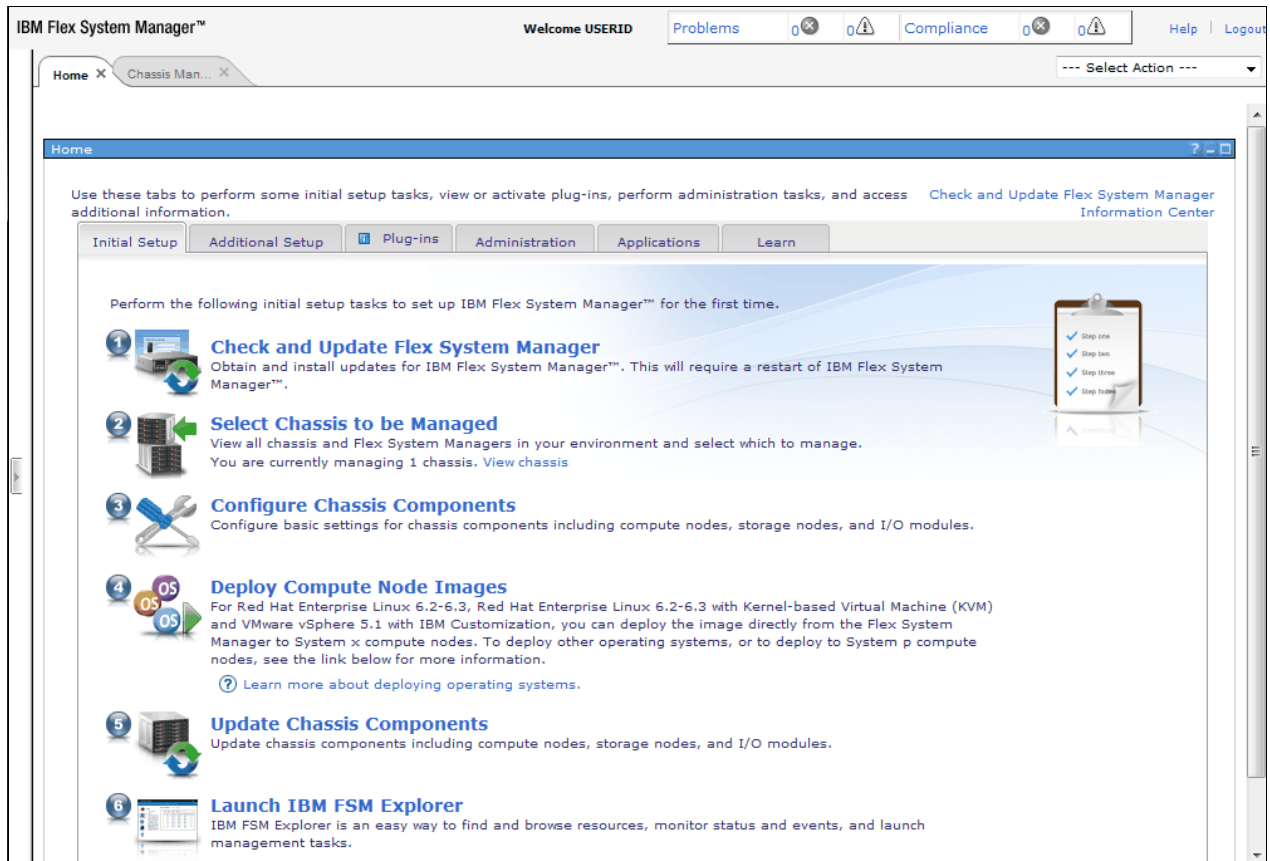


Figure 3-46 FSM - Startup Page

3.3.4 Overview of IBM Flex System Manager and IBM FSM Explorer

The new GUI on IBM Flex System Manager provides two interfaces, which we discuss in the following sections.

IBM Flex System Manager

Next, we give a short overview on how to navigate around in the IBM Flex System Manager.

The home page, shown in Figure 3-46, has six tabs that help you with setting up and administering the system:

Home tab	Initial setup of the system
Additional setup tab	Helps to complete initial setup and prepare IBM Flex System Manager™ for production
Plug-ins tab	Different functions added to the FSM, such as Storage and Network control
Administration tab	A collection of links to ease up configuration and administration
Applications tab	Additional applications added to FSM for more functions, such as IBM Fabric Manager
Learn tab	Links to online education videos regarding the usage of IBM Flex System

The Chassis manager page has multiple functions ,as shown in Figure 3-47.

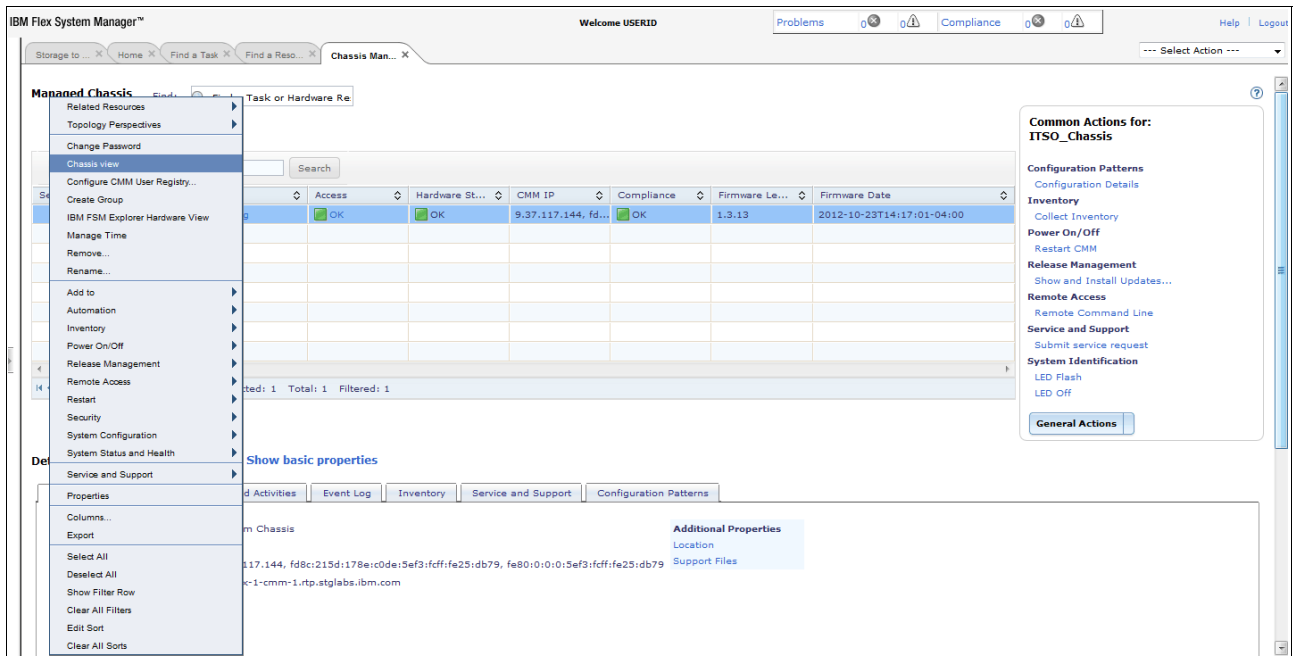


Figure 3-47 FSM chassis manager tab

Figure 3-47 shows the FSM chassis manager. This window offers the following functions:

- ▶ Clicking the chassis name opens IBM FSM Explorer, covered later in this section.
- ▶ The global search tab helps in finding tasks and hardware resources.
- ▶ The action tab gives you advanced configuration options of the chassis.
- ▶ The details section gives you detailed information regarding the chassis, such as eventlog and inventory.

In the FSM GUI, there is a small slide-out tab in the center-left of the interface. Clicking it opens a menu, which shows a comprehensive list of functions that FSM offers. Figure 3-47 shows this FSM slide-out menu.

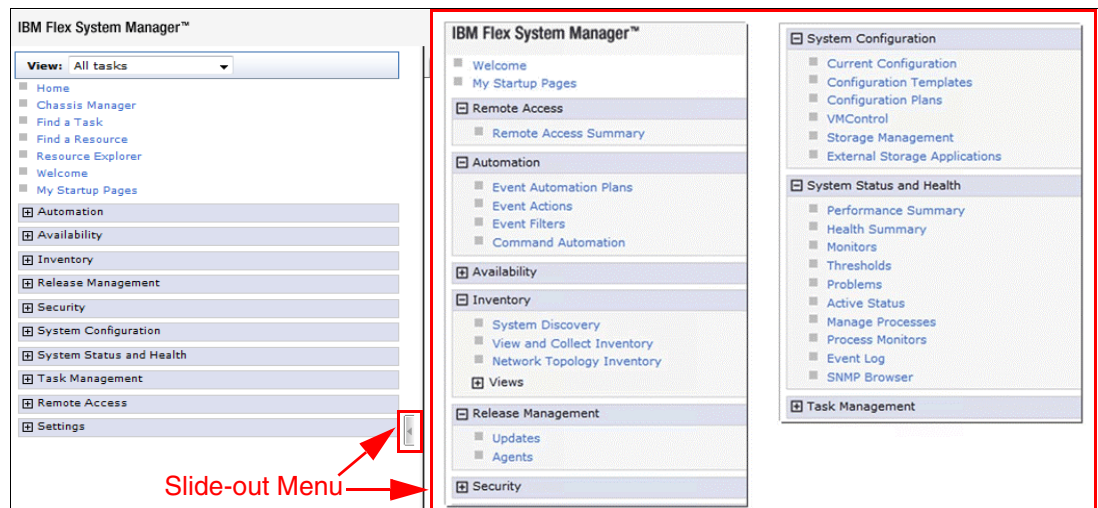


Figure 3-48 FSM - Slide-out menu

As seen in Figure 3-48 on page 118, there are numerous enhanced features and functions that can be easily preformed in the user-friendly FSM user-interface. For detailed information regarding the IBM Flex System Manager web interface and management options, see the *IBM Flex System Manager Systems Management Guide*, located at the following website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.director.main.helps.doc/fsm_1_fqr0_bk_sysm_mgmt.pdf

IBM FSM Explorer

This section gives you a short overview on how to navigate around in the IBM FSM Explorer.

IBM FSM Explorer has a horizontal pull-down menu, and a vertical menu. The horizontal menu is divided into two menus, one in the top black frame, showing health and job status, and the second in the blue frame, which guides you to different tasks and hardware recourses.

To get to IBM FSM Explorer, click **Launch IBM FSM Explorer** in the home tab of IBM Flex System Manager. It will open IBM FSM Explorer in a new tab. Figure 3-49 shows the IBM FSM Explorer home page.

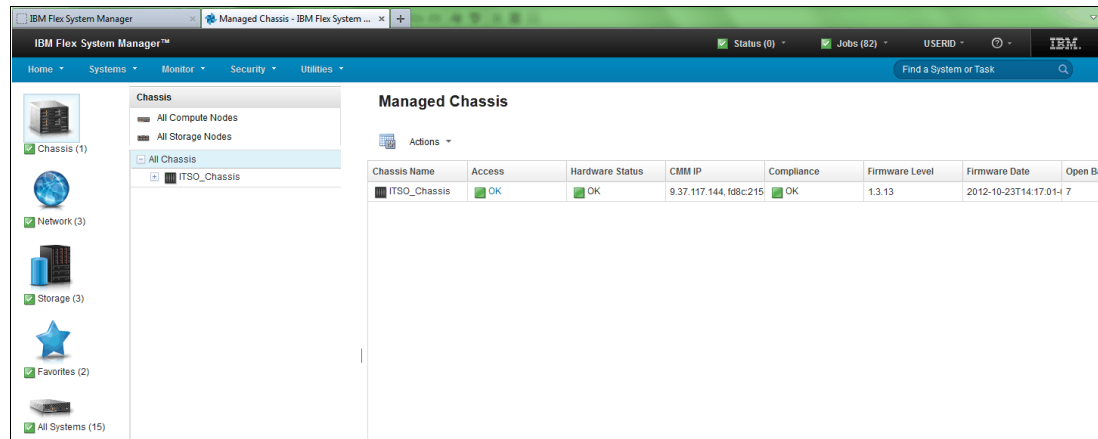


Figure 3-49 IBM FSM Explorer home - tabs

When working in IBM FSM Explorer, it is important to know that some tasks will be opened in IBM FSM Explorer, and other tasks will be opened in IBM Flex System Manager. When choosing tasks in the horizontal menu, it will inform you if the selected task is opened in IBM Flex System manager. Figure 3-50 shows how to identify where a chosen task will open.

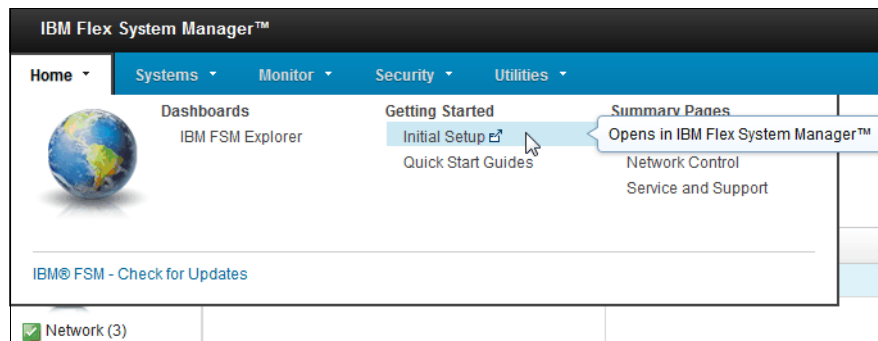


Figure 3-50 IBM FSM Explorer - Tasks

IBM FSM Explorer - Important information: If you choose a task from the chassis hardware map, or any other menu than the horizontal menu, you will not be informed if it is opened in IBM FSM Explorer or in IBM Flex System Manager.

Horizontal menu

The horizontal menu is divided into five pull-down menus, which are covered in this section.

The Home menu is divided into three groups:

- ▶ Dashboards takes you back to the main page.
- ▶ Getting Started gives you shortcuts to initial setup of the FSM and Quick Start Guides in the InfoCenter.
- ▶ Summary Pages gives you shortcuts to Storage control, Network control, and the Service and Support pages in IBM Flex System Manager.

Figure 3-51 shows the Home menu.

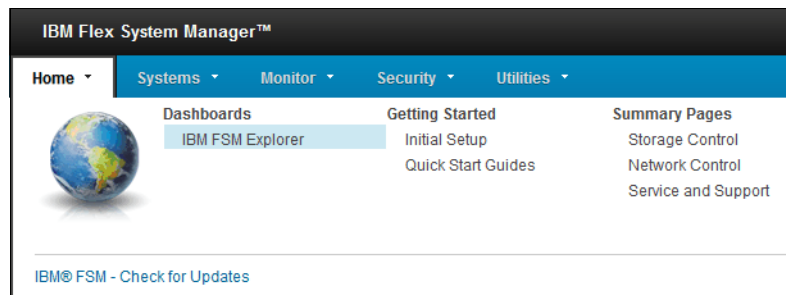


Figure 3-51 IBM FSM Explorer - Home menu

The Systems menu is divided into three groups:

- ▶ Systems gives you shortcuts to all groups defined, the OS group, and Power systems management. Power systems management shows all power systems recourses.
- ▶ Configuration takes you to the configuration patterns.
- ▶ Firmware takes you to Update Manager in the IBM Flex System Manager.

Figure 3-52 shows the Systems menu.

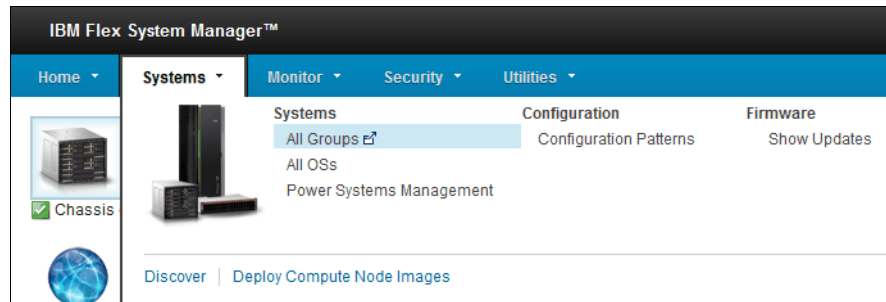


Figure 3-52 IBM FSM Explorer - Systems menu

The Monitor menu is divided into five groups:

- ▶ Status gives you a shortcut to the Active Status page, showing open alerts, Event Log page, and Service Problems page.
- ▶ Monitoring gives you shortcuts to the Monitors page and Threshold page, which allows you to create thresholds and select which systems these thresholds should be used on for monitoring.
- ▶ Compliance gives you shortcuts to Update Manager (Monitored Systems) and Compliance policy page (Policies).
- ▶ Automation gives you shortcuts to the Event Automation Plans page, Event Actions page, Event Filters page, and Command Automation page.
- ▶ Jobs gives you a shortcut to the Active and Scheduled Jobs page.

Figure 3-53 shows the Monitor menu.

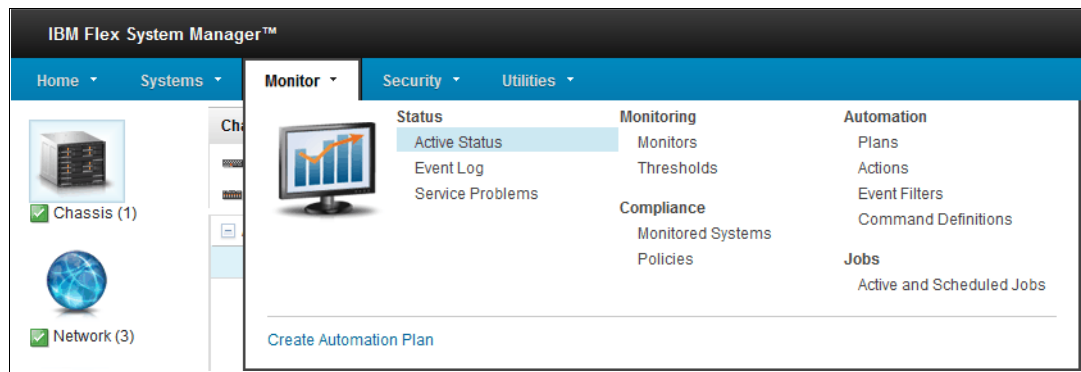


Figure 3-53 IBM FSM Explorer - Monitor menu

Security menu is divided into two groups:

- ▶ Permissions gives you shortcuts to Users and Groups, Roles, and Credentials, which are used to give access and rights.
- ▶ Access takes you to Systems With No Access, which is recourses, which they have not been granted access to.

Figure 3-54 shows the Security menu.



Figure 3-54 IBM FSM Explorer - Security menu

The Utilities menu is divided into three groups:

- ▶ Settings gives you links to the Auto-Rename page, which allows you to rename resources; the Service and Support page, which allows you to configure the system to automatically report errors to IBM; and Server Auditing, which takes you to the Credentials page.
- ▶ Administration gives you links to Administration, Plug-ins, and Applications tab in the IBM Flex System Manager Home page.
- ▶ User Preferences gives you links to FSM (Management Server) and Network Management.

Figure 3-55 shows the Utilities menu.

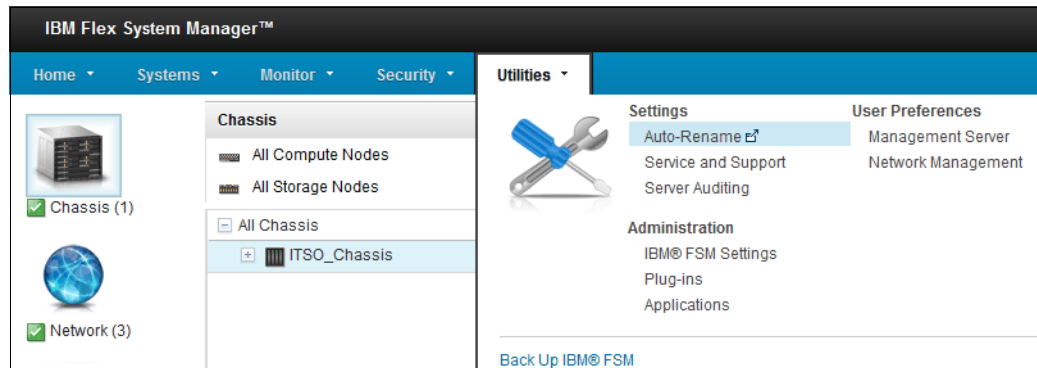


Figure 3-55 IBM FSM Explorer - Utilities menu

Vertical menu

The vertical menu has five images, which open different pages. These pages are covered in the following section.

Figure 3-56 shows the vertical menu.

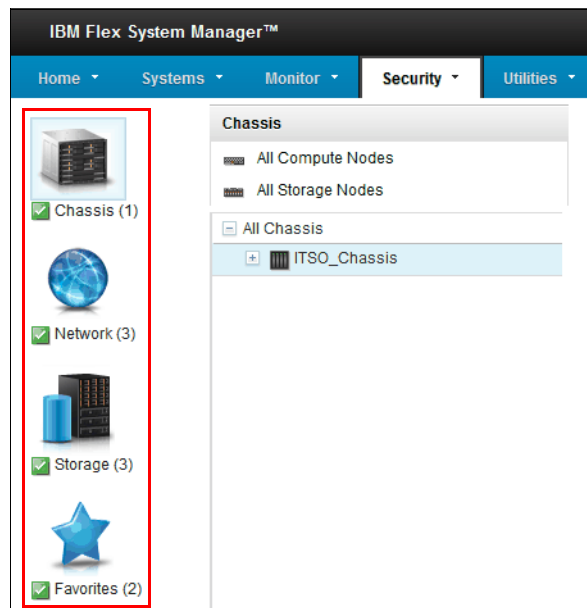


Figure 3-56 IBM FSM Explorer - Vertical menu

The Chassis page gives you an overview of what is installed in one or more chassis, as shown in Figure 3-57.

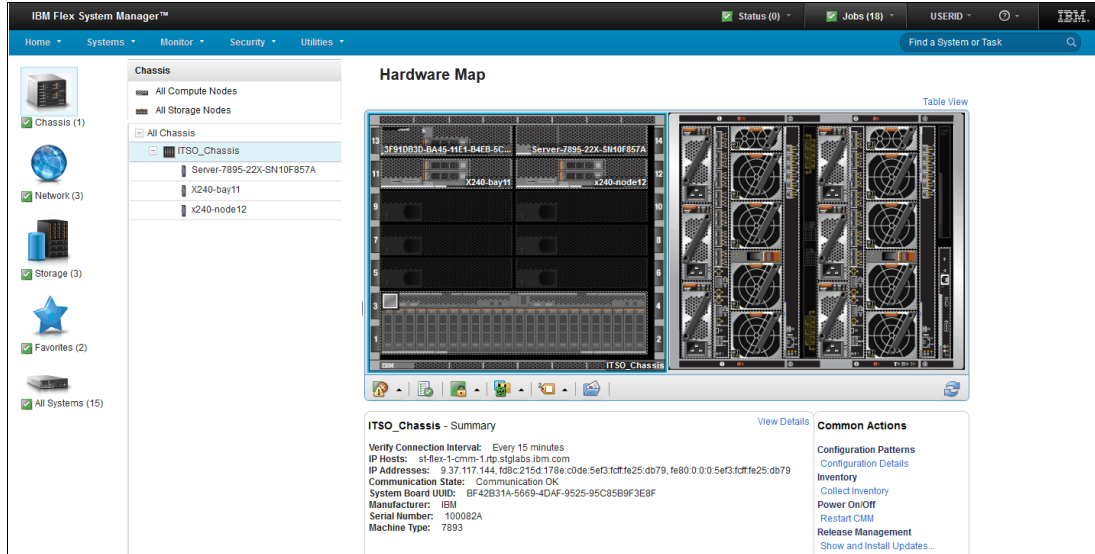


Figure 3-57 IBM FSM Explorer

To get to the chassis map, click the chassis in the vertical menu in Figure 3-56 on page 122, then click the chassis under the chassis list as shown in Figure 3-58.

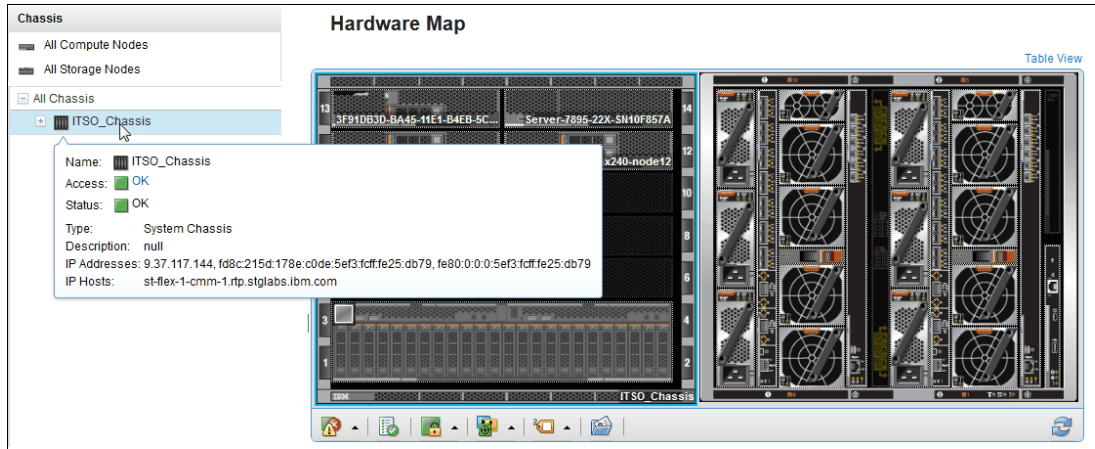


Figure 3-58 IBM FSM Explorer home tab - chassis map

The chassis view shows a map of the front and back with a graphical representation of the components, as shown in Figure 3-59.

The screenshot displays the 'Hardware Map' interface in IBM FSM Explorer. It features two main graphical views: a front view on the left and a back view on the right. The front view shows a server chassis with numbered slots (1-14) and various components. The back view shows the internal components, including fans and drive bays. Below the graphical views is a toolbar with icons for navigation and actions. At the bottom, there is a summary panel for 'ITSO_Chassis' with the following details:

ITSO_Chassis - Summary		View Details	Common Actions
Verify Connection Interval:	Every 15 minutes		
IP Hosts:	st-flex-1-cmm-1.rtp.stglabs.ibm.com		
IP Addresses:	9.37.117.144, fd8c:215d:178e:c0de:5ef3:fcff:fe25:db79, fe80:0:0:0:5ef3:fcff:fe25:db79		
Communication State:	Communication OK		
System Board UUID:	BF42B31A-5669-4DAF-9525-95C85B9F3E8F		
Manufacturer:	IBM		
Serial Number:	100082A		
Machine Type:	7893		
			Configuration Patterns Configuration Details Inventory Collect Inventory Power On/Off Restart CMM Release Management Show and Install Updates...

Figure 3-59 IBM FSM Explorer chassis map

To get detailed information about the chassis and chassis components, click **View Details** in the summary field of Figure 3-59.

As shown in Figure 3-60, *Common Actions* for chassis are shown below the Hardware map at the left side. The detailed information regarding the chassis is shown below the Hardware map at the right side.

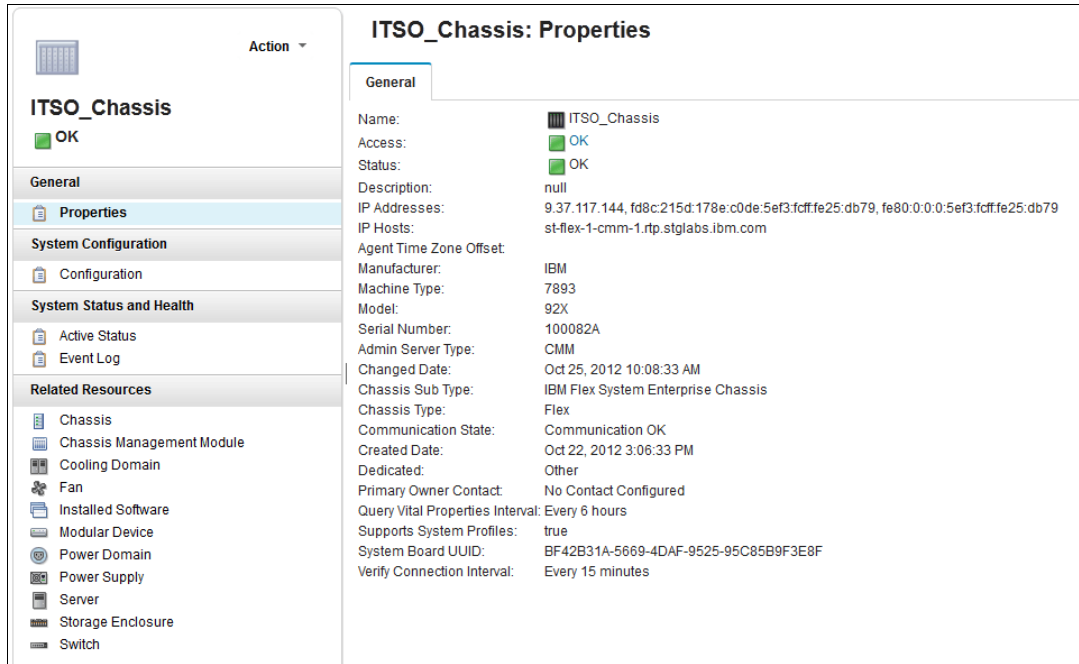


Figure 3-60 IBM FSM Explorer chassis details

A drop-down menu shows the actions that can be performed at the Chassis level, as shown in Figure 3-61.

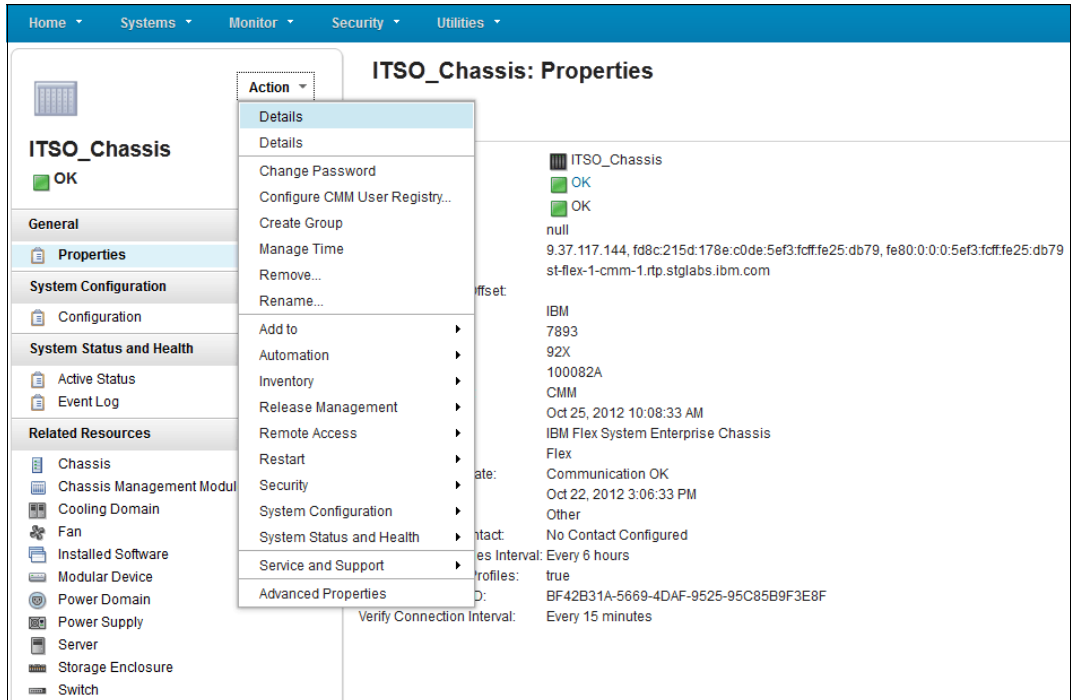


Figure 3-61 IBM FSM Explorer - Component Actions

The Network page gives you an overview of VLANs, Subnets, and network switches, as shown in Figure 3-62.

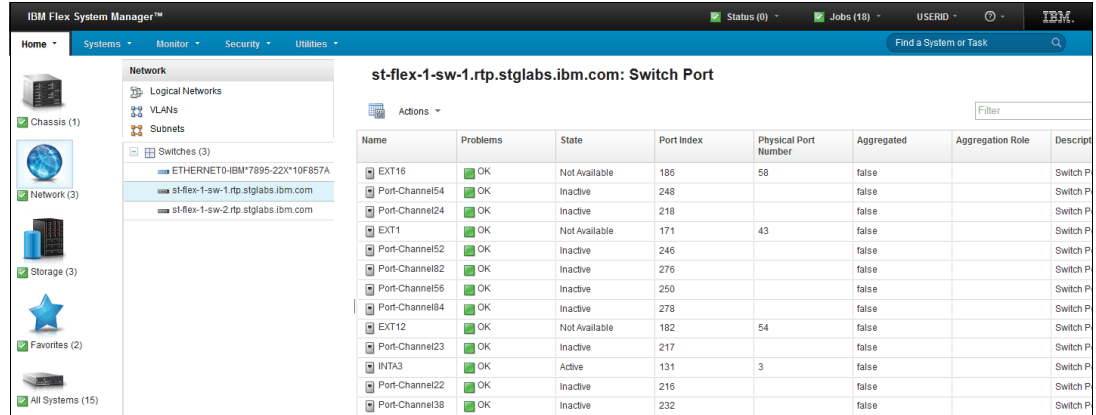


Figure 3-62 IBM FSM Explorer - Network page

The Storage page gives you an overview of all storage systems, Fabrics, and Storage pools, as shown in Figure 3-63.

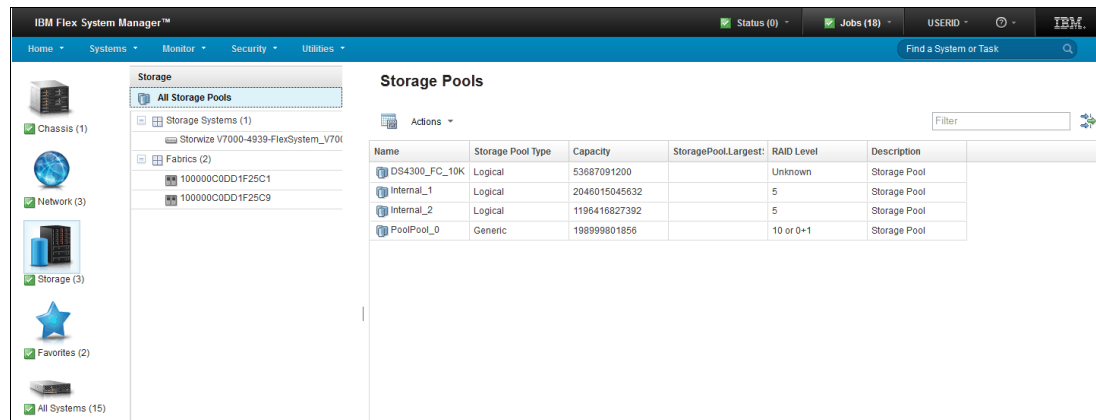


Figure 3-63 IBM FSM Explorer - Storage page

Right-clicking an IBM Flex System V7000 Storage Node gives you the possibility to update firmware as shown in Figure 3-64.

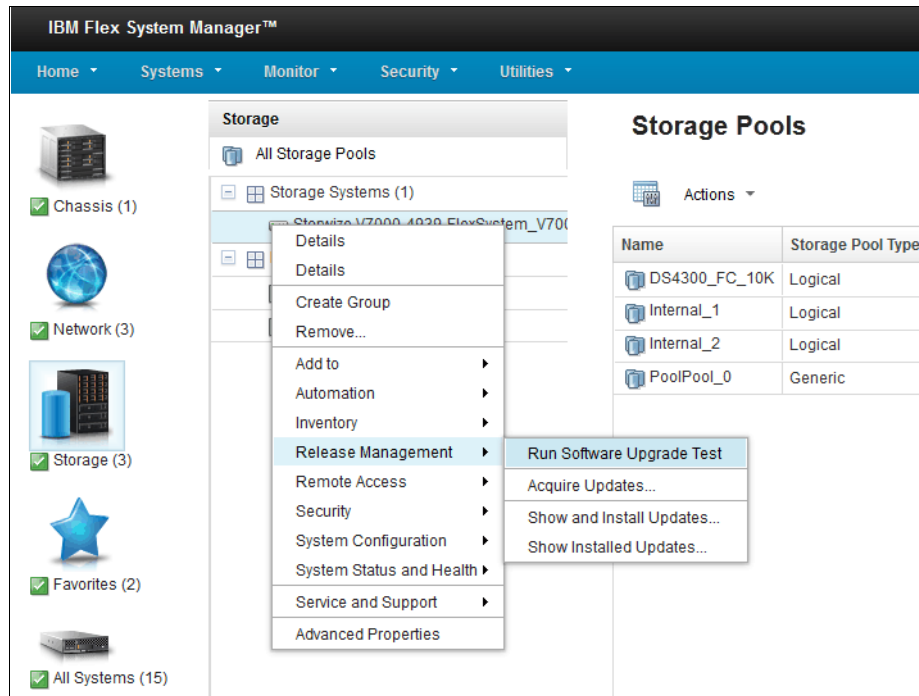


Figure 3-64 IBM FSM Explorer - Storage page - Firmware update

From the menu shown in Figure 3-64, you can also launch the IBM Storwize V7000 Element manager, as seen in Figure 3-65.

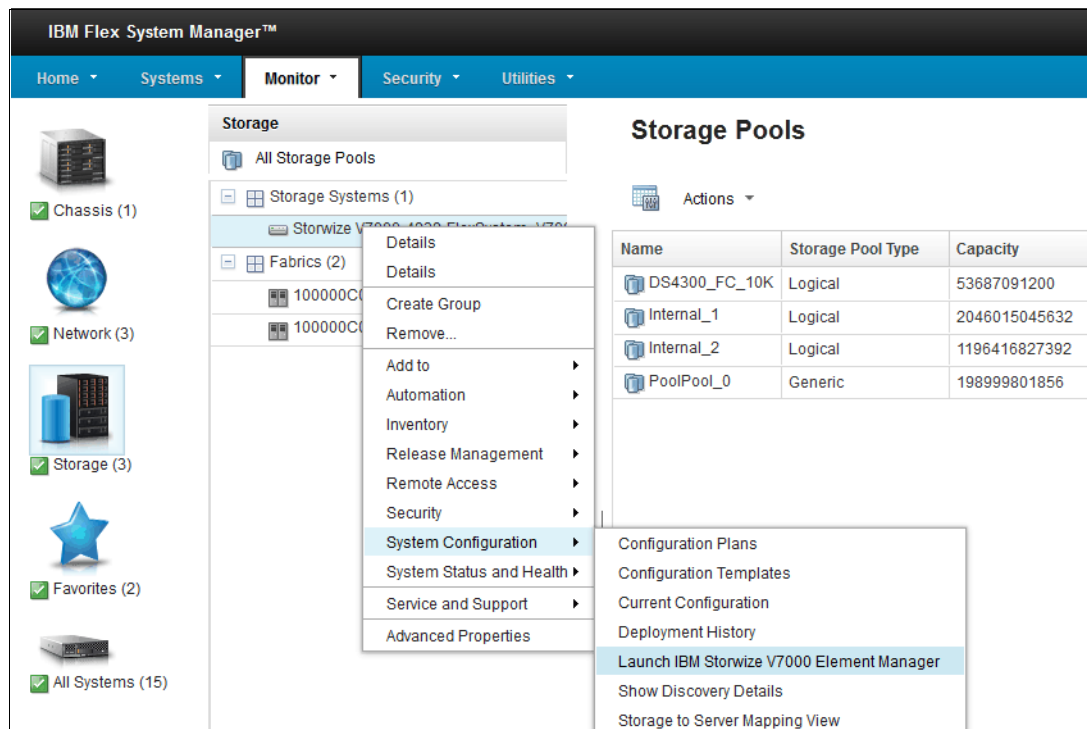


Figure 3-65 IBM FSM Explorer - Storage page - launch IBM Storwize V7000 Element manager

Choosing the option to Launch the IBM Storwize V7000 Element Manager will open the IBM Flex System V7000 Storage Management login page, as shown in Figure 3-66.



Figure 3-66 IBM Flex System Storage Management - Login page

In the Fabrics section, clicking one of the numbers (Last 12 numbers is the MAC address of the switch) will show you all the switches, as shown in Figure 3-67.



Figure 3-67 IBM FSM Explorer - Storage page - Switch overview

Double-clicking one of the switches leads you to detailed information about the switch, as shown in Figure 3-68.

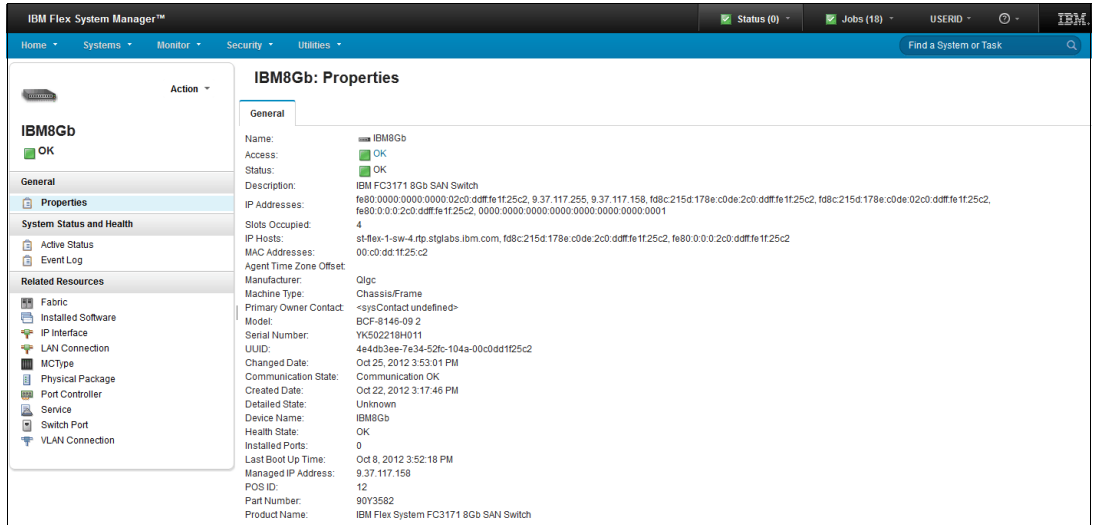


Figure 3-68 IBM FSM Explorer - Storage page - Switch details

From the switch overview page, right-click one of the switches, and select **Remote access** → **Launch Web Browser**. It will open the switch web GUI login prompt, from where the switch can be configured. Figure 3-69 shows how to open the switch web GUI.

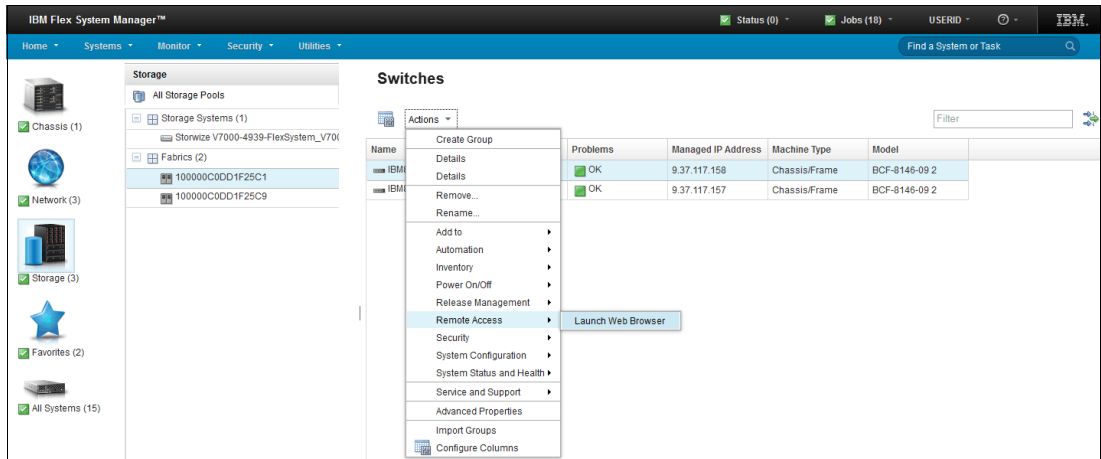


Figure 3-69 IBM FSM Explorer - Storage page - Switch GUI

Figure 3-70 shows the FC switch login page.

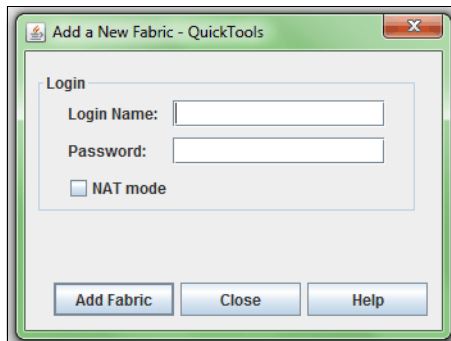


Figure 3-70 FC Switch login

3.3.5 Accessing I/O modules using FSM

In an IBM Flex system, node to node communication happens within the Chassis.

Network Control

Network Control is a plug-in that enables utilization of integrated network management capabilities of an FSM. It has the following key features:

- ▶ Logical network management allows management of port profiles, VLANs, ACLs, and QoS in virtualized, live-migration environments.
- ▶ It uses 802.1Qbg standards in integrated switches and PowerVM, KVM, and IBM “Osiris” vSwitch for VMware.
- ▶ Optional Fabric Management extends QoS Management by providing advanced monitoring, VM priority, and rate limiting.
- ▶ Network monitoring at a glance is provided by network topology perspectives with the ability to see the components affected by network outages.
- ▶ It enables end-to-end network and virtualization management.

To access Network Control, from IBM FSM Explorer, select **Home menu** → **Network Control**, as shown in Figure 3-71.

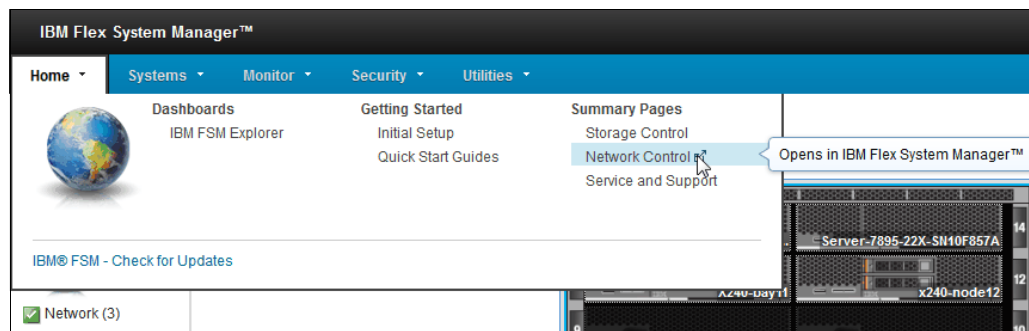


Figure 3-71 IBM FSM Explorer - Network control

This action will open up the Network Control interface, which shows management and configuration options for the IBM Flex System I/O Modules, as shown in Figure 3-72.

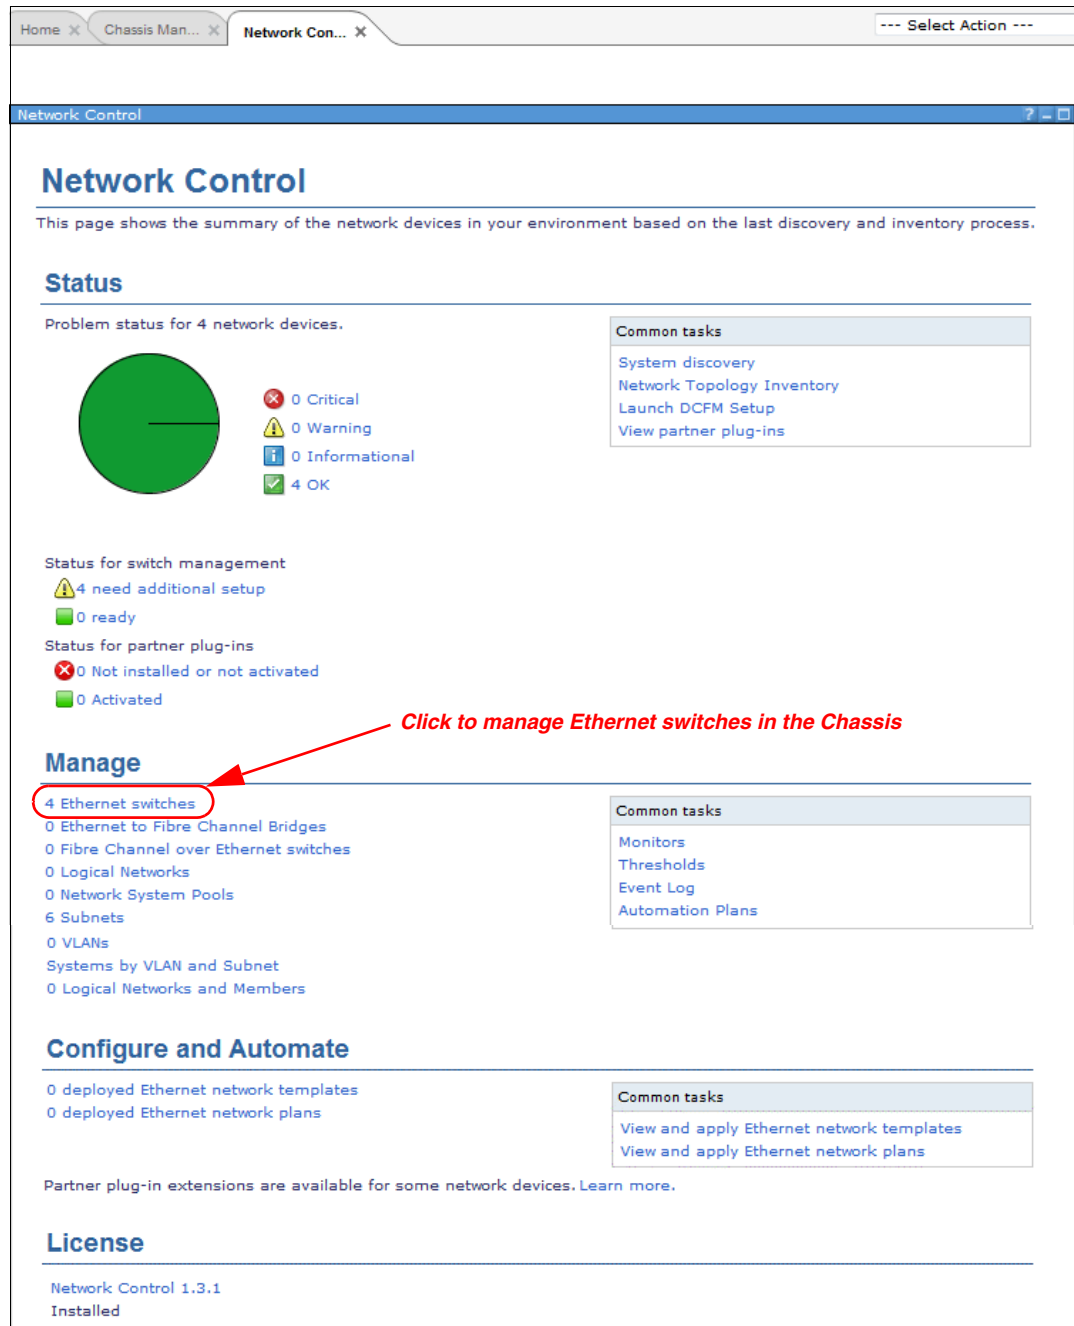


Figure 3-72 Network Control Plug-in

Selecting **Ethernet Switches** in Network Control interface opens a view of all available Ethernet Switches along with information regarding state, access, IP and type, as shown in Figure 3-73.

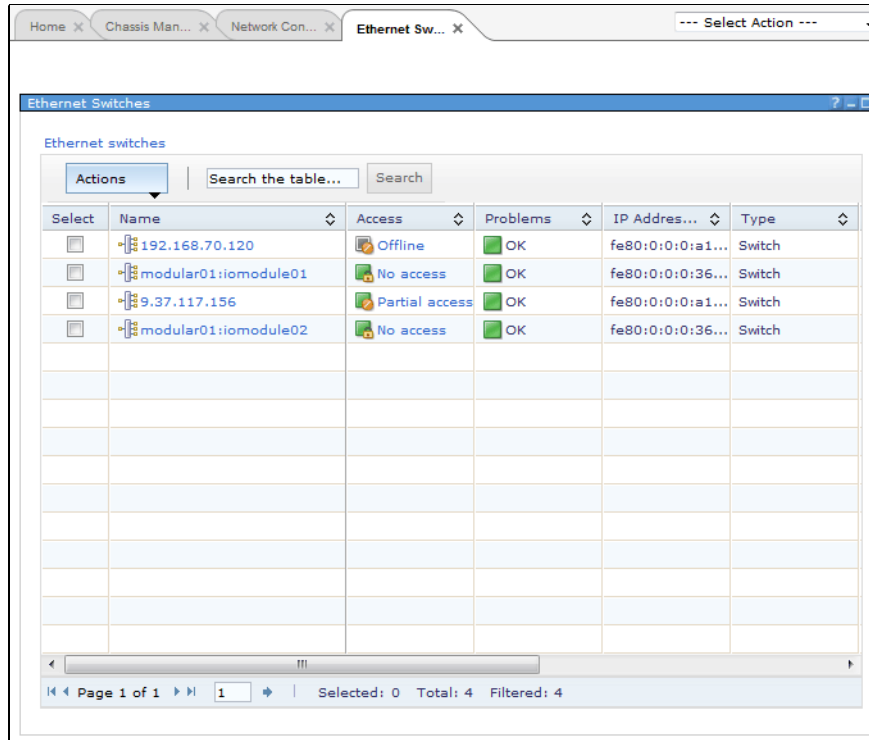


Figure 3-73 Viewing Ethernet Switches in FSM

Select a switch by clicking its given *name* to go to the detailed management and configuration options for it, as shown in Figure 3-74.

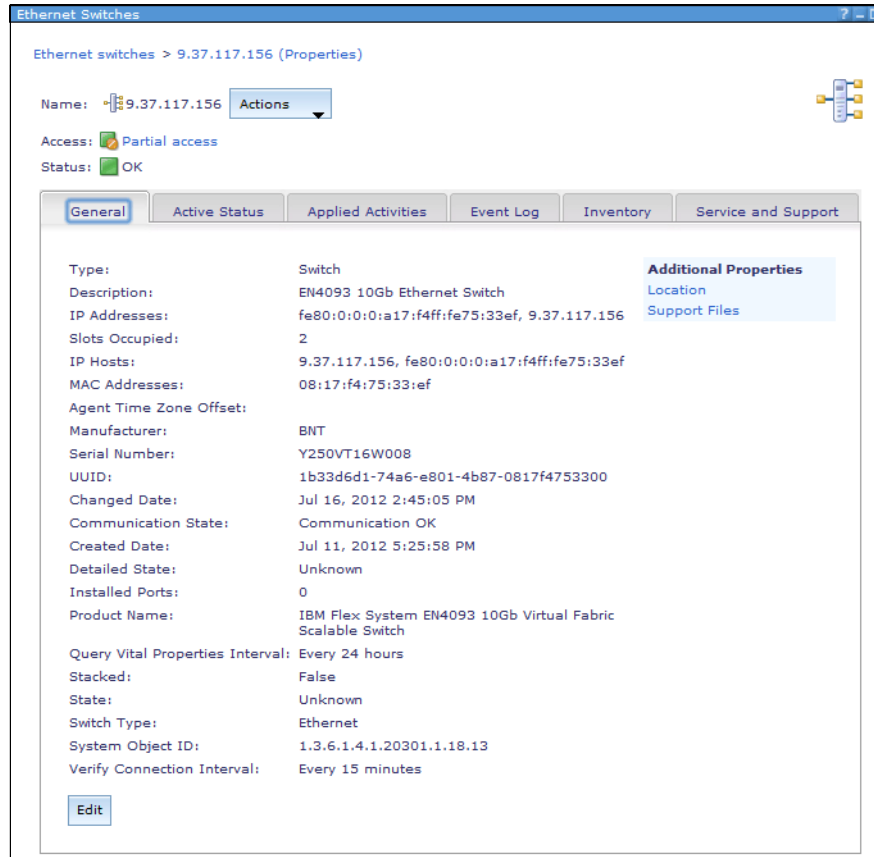


Figure 3-74 Details regarding selected ethernet switch

Several tabs in this view illustrate enhanced monitoring and logging capabilities of the FSM. Click the **Inventory** tab to view and gather Inventory for the selected switch module, as shown in Figure 3-75.

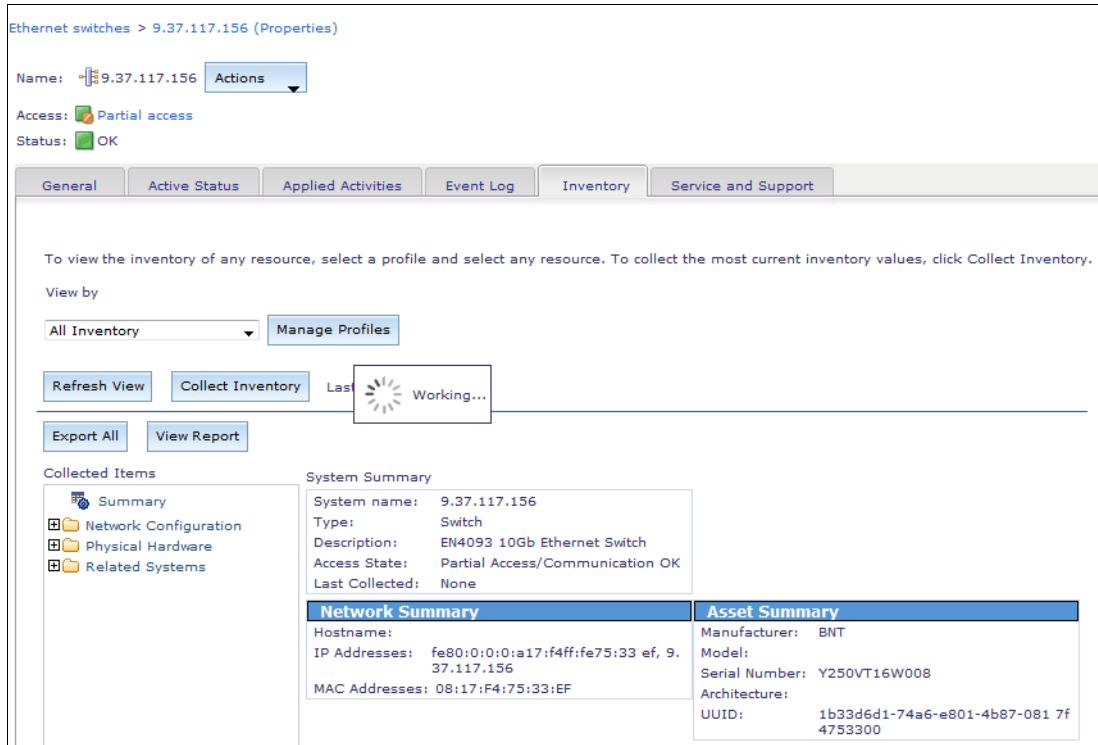


Figure 3-75 Gathering Inventory information

References

For further details regarding the management and configuration of I/O Modules using IBM Flex System Manager, visit the following website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.8731.doc/managing_network_resources.html

3.3.6 Data collection using FSM

You can use IBM Flex System Manager to collect and submit support files for a managed resource. Support files can contain detailed system information used to help diagnose a serviceable hardware problem, dump files collected from a managed resource, event logs, and more.

Automatic versus manual data collection

By default, a plug-in FSM tool called Service and Support Manager automatically collects additional data associated with a serviceable hardware problem, and stores it as a support file. However, you can also collect support file data manually, even when a serviceable hardware problem has not occurred.

Service and Support Manager (SSM)

IBM Service and Support Manager monitors, tracks, and captures system hardware errors and service information, and reports serviceable problems directly to IBM Support using the IBM Electronic Service Agent (ESA) tool. The IBM Flex System Manager will automatically monitor for any serviceable problems on your managed hardware and automatically report problems to IBM service.

IBM Service and Support Manager offers enhanced reporting abilities. It opens service requests for failed hardware using embedded Electronic Service Agent (ESA). It also collects, transmits, and manages extended error data including logs and diagnostics. Another added advantage of SSM is transmission of hardware and software inventory data to IBM using ESA, which assists in troubleshooting the problem as well as monitor the dependencies.

Tip: If you do not set up the Electronic Service Agent (ESA) then problems will be collected in the FSM management appliance but will not be reported to IBM, however, you can always use the FSM itself to monitor and track your serviceable problems.

To access the Service and Support Manager (SSM), click **Home menu** → **Service and Support**, as shown in Figure 3-76.

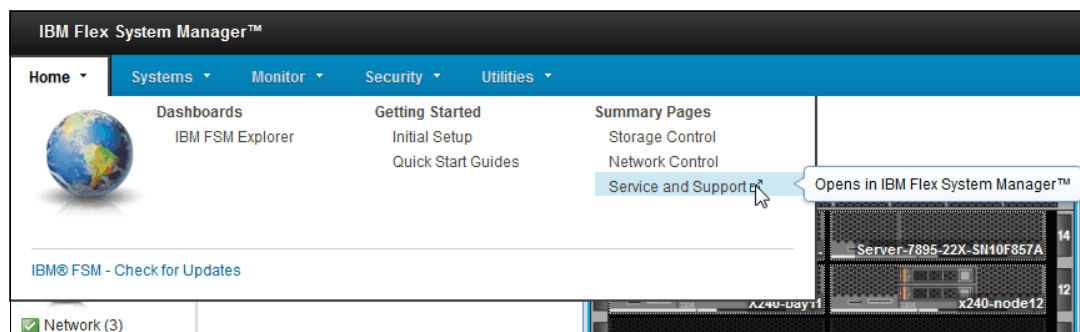


Figure 3-76 Opening Service and Support Manager (SSM)

The SSM interface displays, which includes a comprehensive set of Service Management tools in one tabular view, as shown in Figure 3-77.

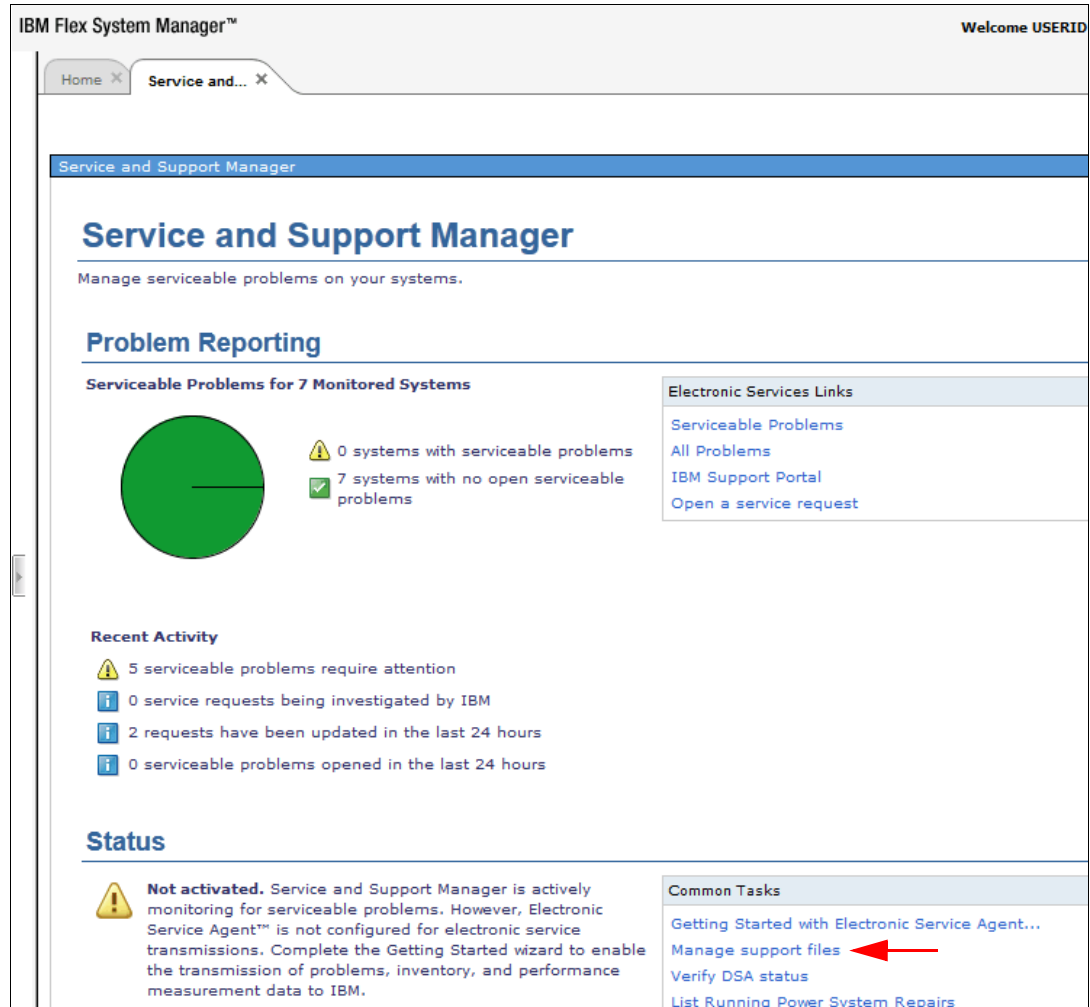


Figure 3-77 SSM default view

To collect logs from the IBM Flex System V7000 Storage Node using IBM Flex System Manager, click **Manage support files** as shown in Figure 3-77.

The “Manage support files” page displays, and by choosing **Collect support files**, a new window appears, where you can see logs already collected. as shown in Figure 3-78 .

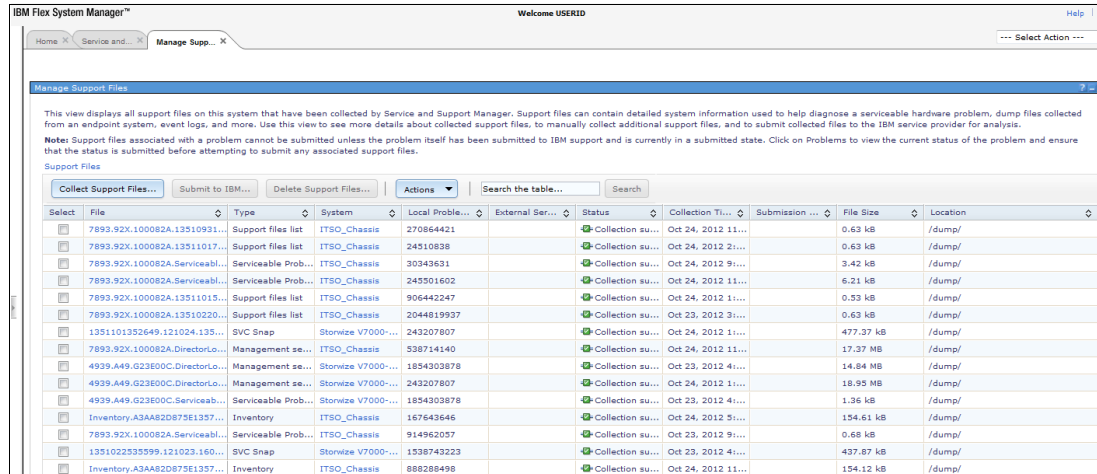


Figure 3-78 IBM FSM - Manage support files

Select one or more files, then click **Actions**, and you have several options to download or transfer the logs:

- ▶ Submit to IBM: (only works if Electronic Service Agent is enabled)
- ▶ Upload to ftp server: Upload to an ftp server on your network
- ▶ Copy to media: USB key, inserted into IBM Flex System Manager Node
- ▶ Download: Download the support file to the computers which are connected to the IBM Flex System Manager

Figure 3-79 shows the options for a support file.

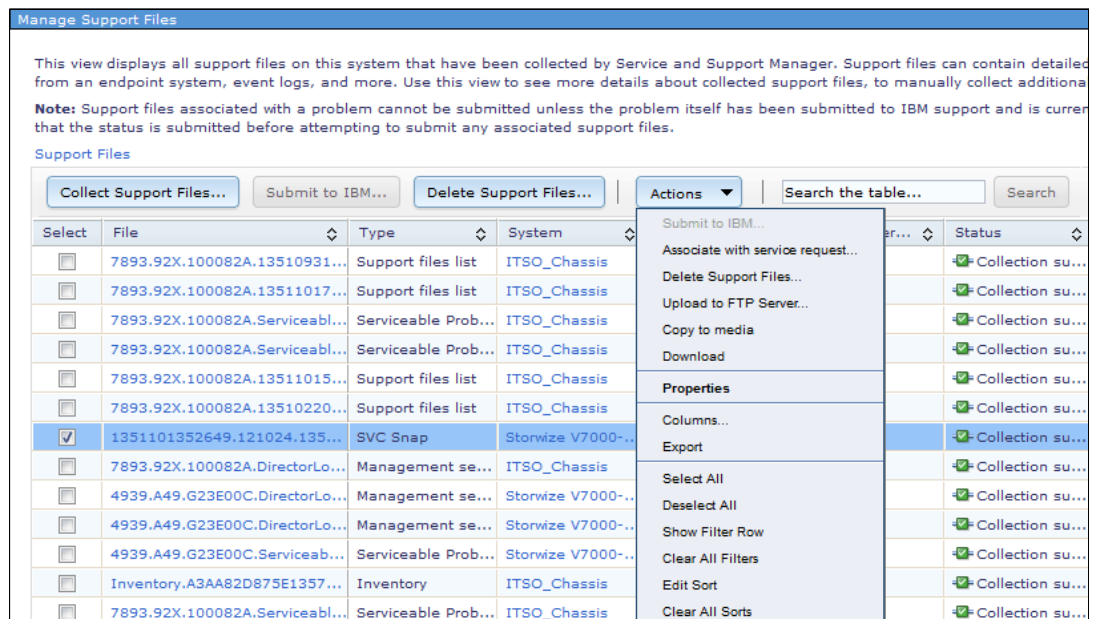


Figure 3-79 IBM FSM - Manage support files download

To collect a new support file, click **Collect Support Files**. A new window displays, where you can choose the device, as shown in Figure 3-80.

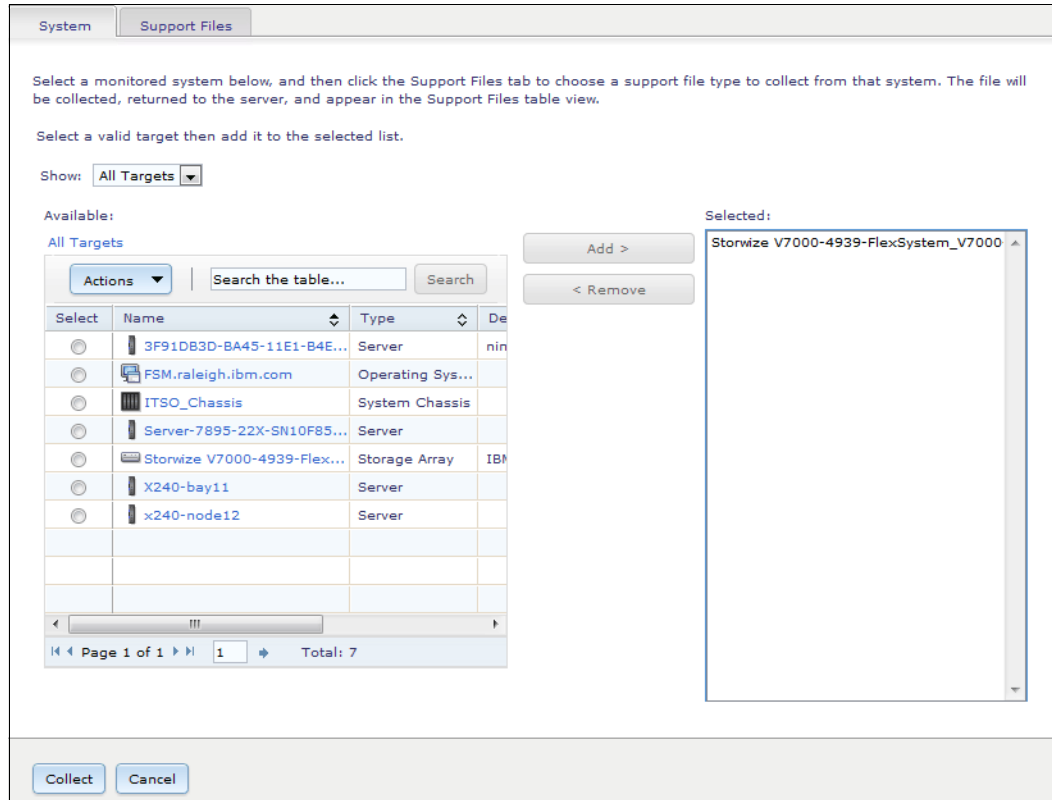


Figure 3-80 IBM FSM - Collect support file - choose device

To select the type of support file you which to collect, click the **Support Files** tab, then select the file and click **Collect** as shown in Figure 3-81.

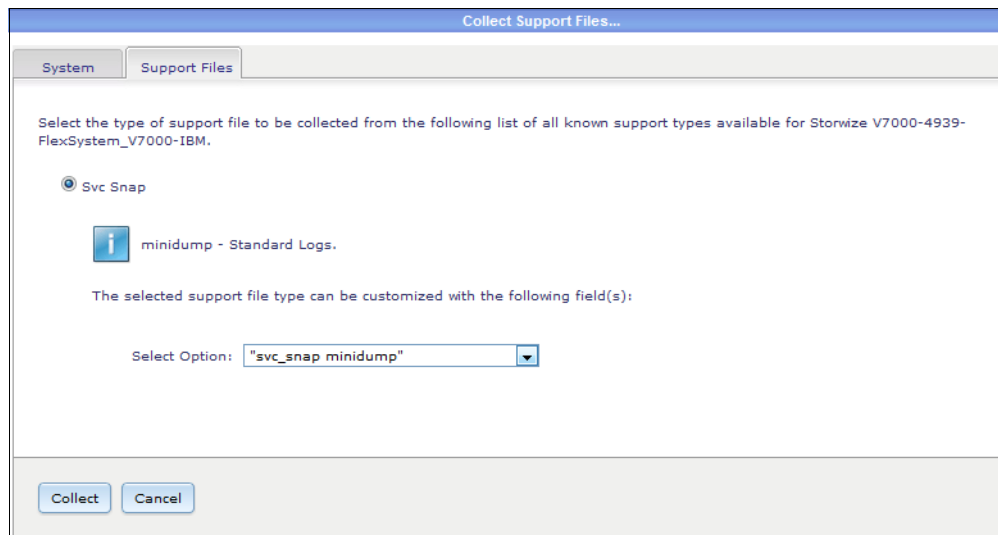


Figure 3-81 IBM FSM - Collect support file - choose file type

When the file has been collected, you can choose to download it, send it to IBM or copy it to a media.

3.3.7 Managing storage using IBM Flex System Manager

Storage management with IBM Flex System Manager involves two software components: Storage Manager and Storage Control. Both are included with the management software:

- ▶ Storage Manager is a standard management software capability that provides basic storage life cycle management (Discovery, Inventory, Health and Alerts).
- ▶ Storage Control is a pre-installed plug-in for the management software that expands storage support to mid-range and high-end storage devices.

The Storage Control feature of IBM Flex System Manager allows management of most IBM mid-range and high-end storage systems. With the Storage Control plug-in, you can manage an expanded set of storage subsystems and Fibre Channel switches.

Storage Control offers the following key advantages:

- ▶ Extends storage management of IBM Systems Director and VMControl to cover most IBM storage systems
- ▶ Provides storage device discovery and coverage in integrated physical and logical topology views
- ▶ Shows relationships between storage and server resources
- ▶ Facilitates single management server deployment
- ▶ Provides ability to view controller and volume status and to set notification alerts
- ▶ Virtualizes for higher storage utilization
- ▶ Balances high performance and cost for mixed workloads
- ▶ Protects data and minimizing downtime
- ▶ Integration with VMControl featuring the following additional functions:
 - Storage provisioning for image creation, deployment, and cloning
 - Ability to manage storage system pool life cycle, take group actions across pool and policy-based storage placement, provisioning, and cloning actions within the pool

Storage Control also extends the management capabilities and device coverage of the VMControl, which is a tool within the IBM Flex System Manager Node responsible for the life cycle management of virtual resources. This collaboration is leveraged in VMControl for tasks such as working with system pools for items like server and storage pool creation, the image library management and storage, and integrated virtualization management across server and storage infrastructure.

For further details regarding storage management features, visit the following website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.director.storage.helps.doc/fqm0_t_sm_configuring_storage.html

Storage management and storage control

You can use Storage Control to monitor the health of storage subsystems, storage pools, storage volumes, and fabric switches. You can run inventory against the Storage Control farm to see updated health status.

You can also use Storage Control to manage several external storage subsystems. These subsystems include the IBM® DS3000, DS4000®, DS5000, families, the IBM DS8000® family (DS8100, DS8300, DS8700, DS8800), the IBM SAN Volume Controller, the IBM Storwize V7000, IBM XIV®, Qlogic 8GB Fibre Channel switch, and the McDATA Sphereon and Intrepid switches.

Depending on the firmware levels of these devices, Storage Control supports native interfaces to the device, which simplifies configuration setup and improves device management reliability. After being discovered, these devices will be listed as being managed by Storage Manager in Flex System Manager.

To open Storage Management, choose **Home** → **Storage Control** in the horizontal menu in IBM FSM Explorer as shown in Figure 3-82.

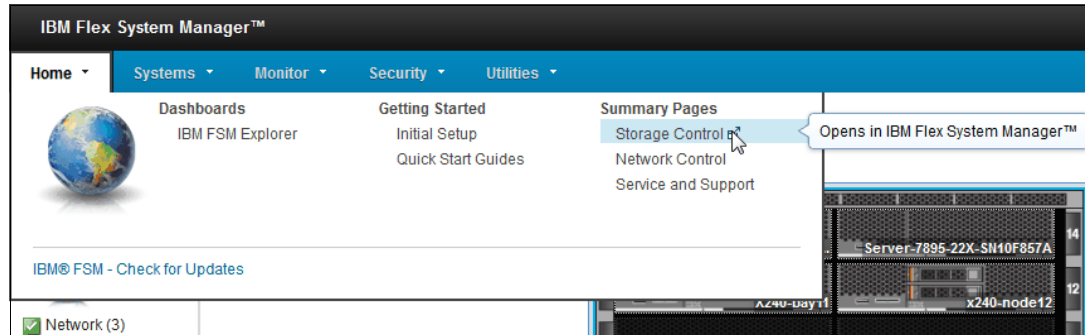


Figure 3-82 IBM FSM Explorer - Storage Control

The Storage Management window opens, showing all storage that is available in the managed IBM Flex Systems, as shown in Figure 3-83.



Figure 3-83 FSM Storage Control

You can also manage storage by going to the **Home tab** → **Additional Setup** → **Manage System Storage**, as shown in Figure 3-84.

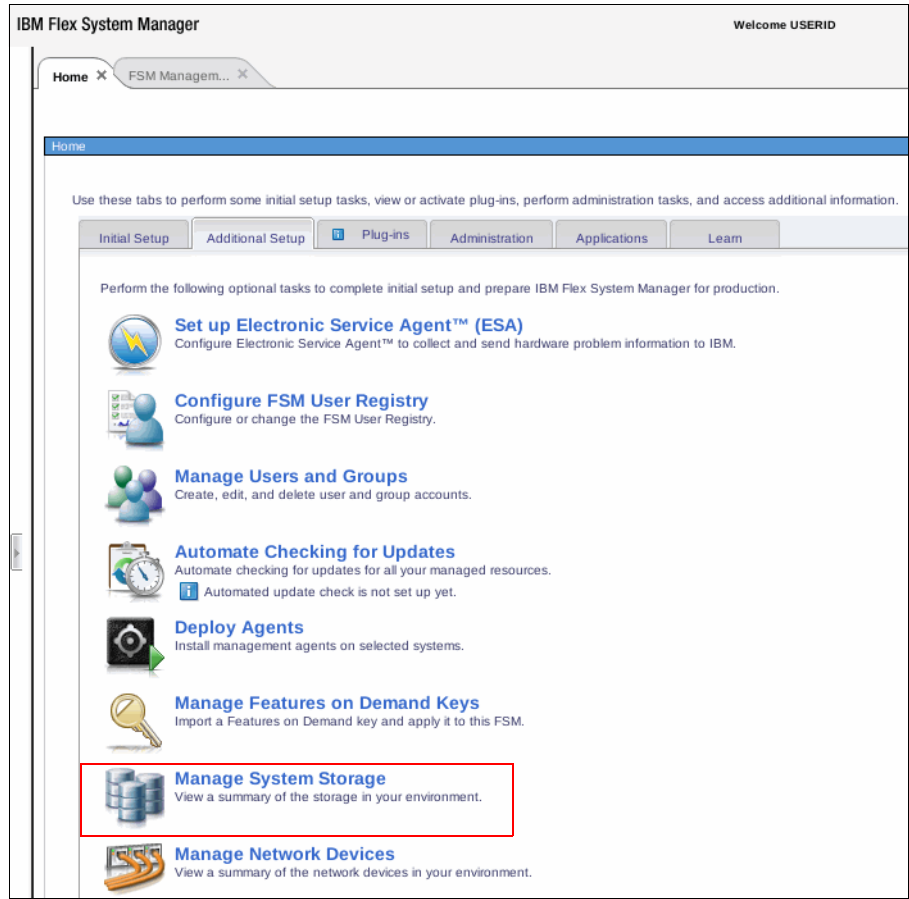


Figure 3-84 FSM - Manage System Storage

To get more information about the storage systems controlled by IBM Flex System Manager and see which actions there can be performed, click **Network Storage** in the storage control window, as shown in Figure 3-85.

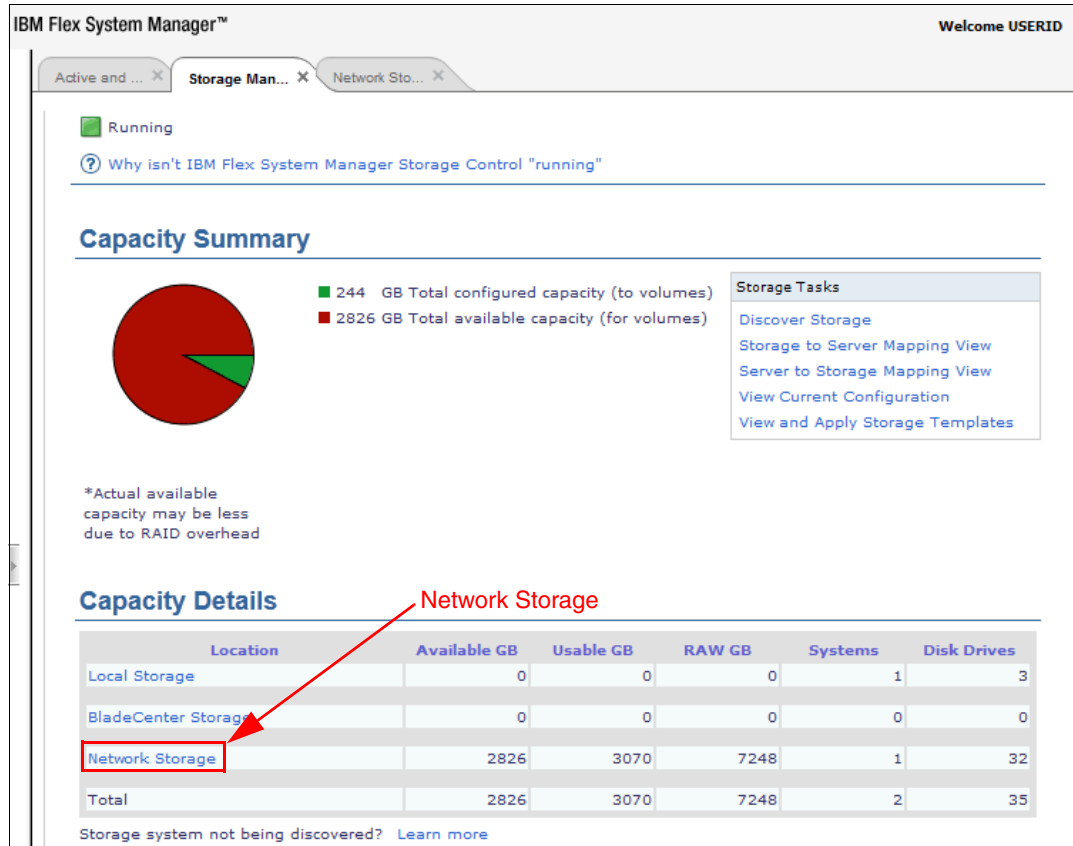


Figure 3-85 FSM - Storage Control - Network storage

This will open a new page showing Network attached storage, as shown in Figure 3-86.

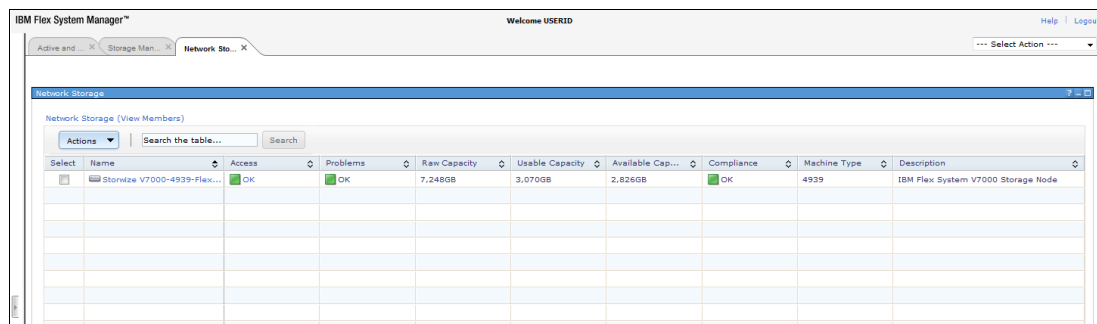


Figure 3-86 FSM - Storage Control - Storage systems

Clicking the storage system name opens the selected storage systems properties page, as shown in Figure 3-87.

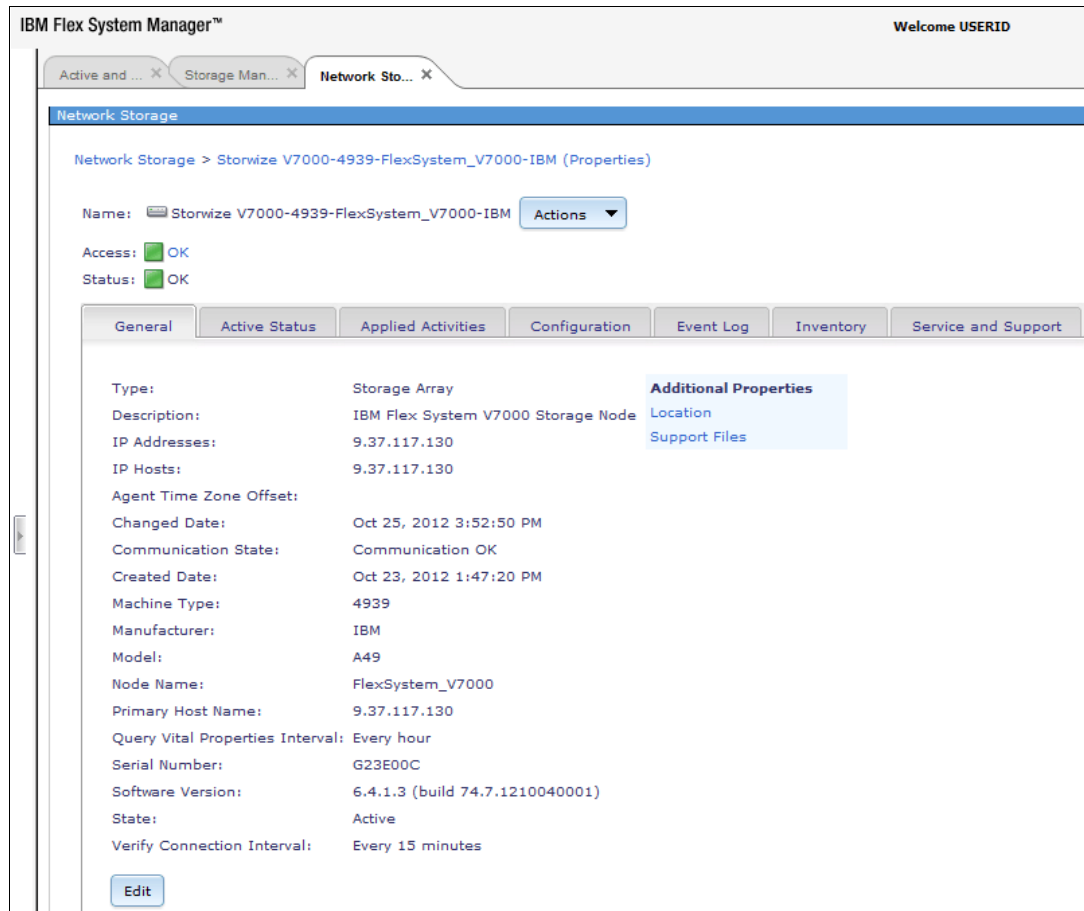


Figure 3-87 FSM - Storage Control - Storage systems

From this window, you are able to see General information, Eventlog, and Inventory of the storage system, by clicking the different tabs.

Choosing the **Actions** pull-down button will allow you to collect inventory, update storage system firmware, and open the IBM Flex System V7000 Storage Node management application, as shown in Figure 3-88.

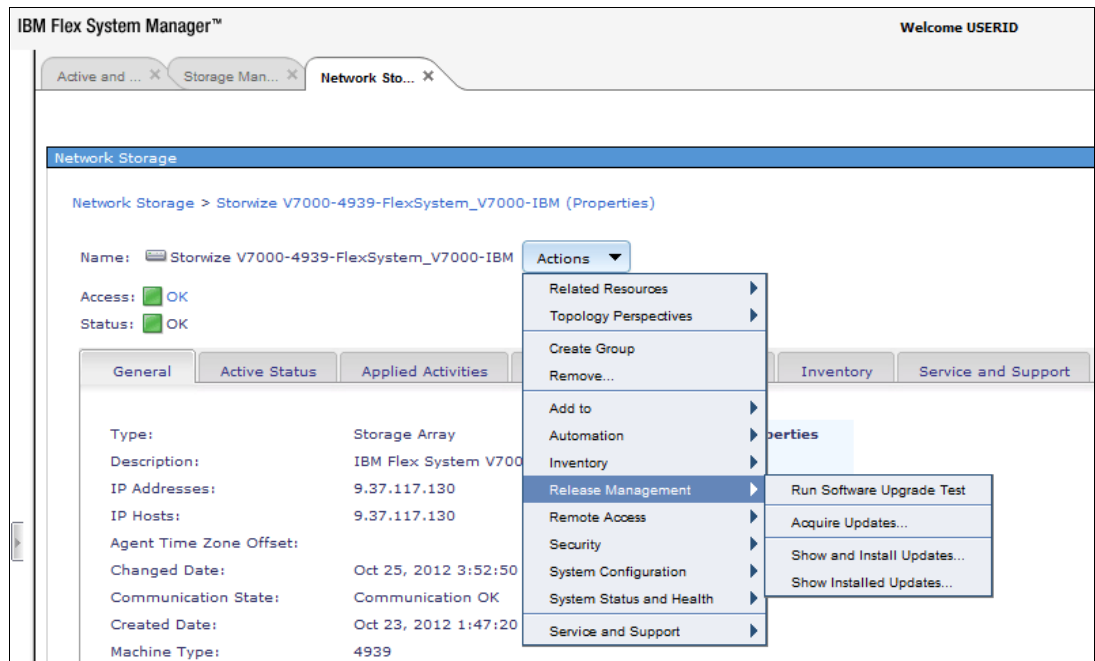


Figure 3-88 FSM - Storage Control - Storage system action menu

Managing external storage systems

To manage external storage subsystems in IBM Flex System Manager, perform the following steps:

1. From the left slide-out tab, select **Resource Explorer**, as shown in Figure 3-89.

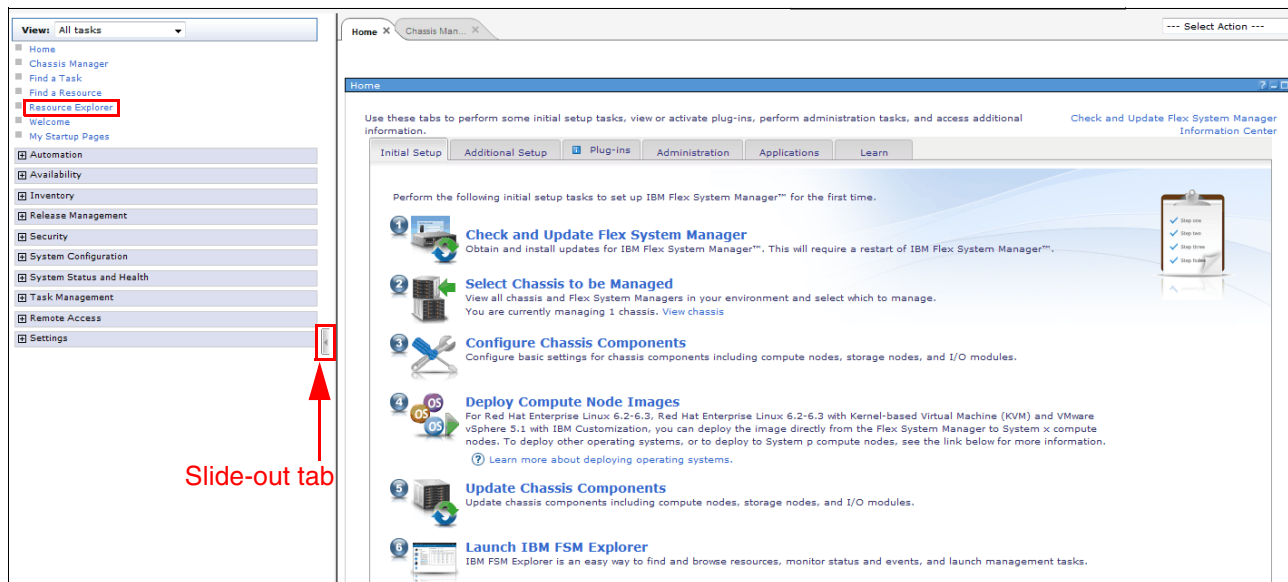


Figure 3-89 Opening Resource Explorer in FSM

2. Resource Explorer shows a snapshot of all resources in the IBM Flex System Manager managed by the IBM Flex System Manager, as shown in Figure 3-90.




Figure 3-90 Resource Explorer

3. Navigate to manage discovered storage, as shown in Figure 3-91.

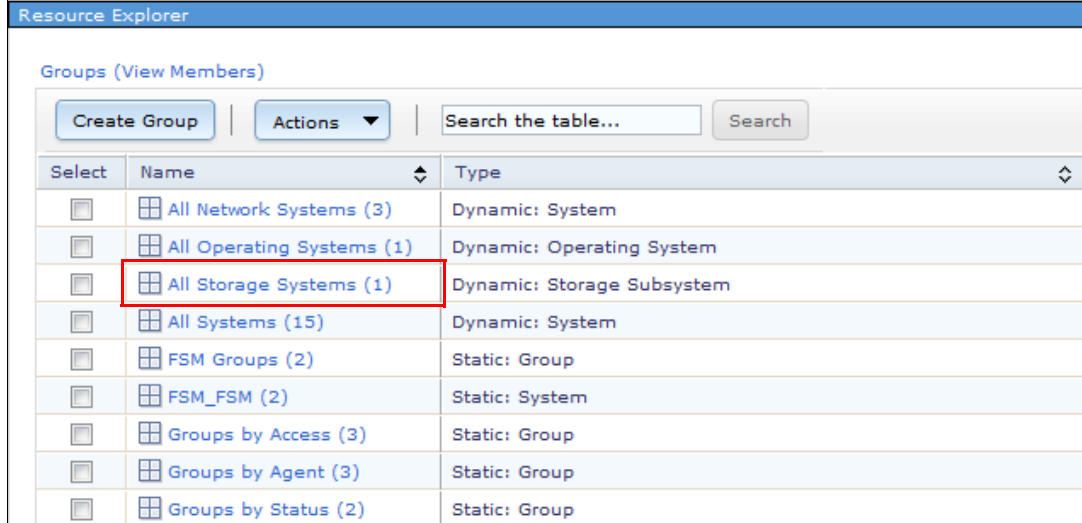


Figure 3-91 Navigate to all storage subsystems

4. Select target storage (for example, IBM Storwize V7000), as shown in Figure 3-92.




Figure 3-92 Selecting target storage

5. Select the tab Inventory to collect the inventory, as shown in Figure 3-93.

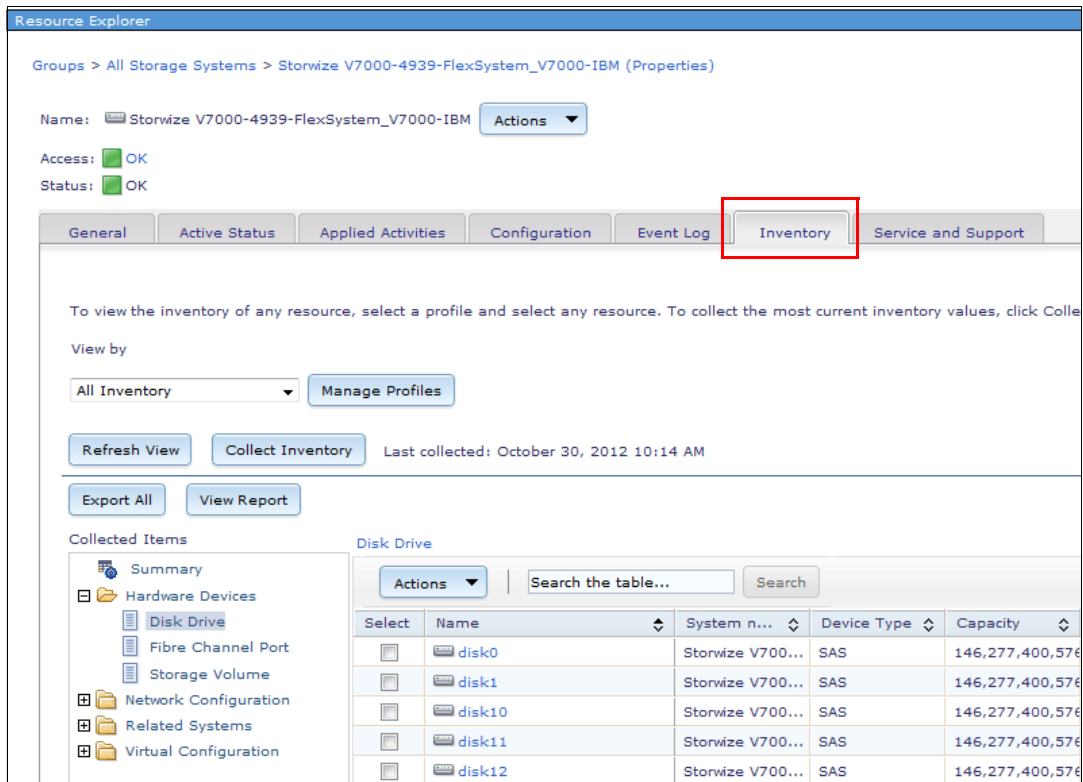


Figure 3-93 Inventory of Storage

6. Use the Task Launch Dialog box that is opened to verify the name of the inventory, and time for running it, as shown in Figure 3-94.

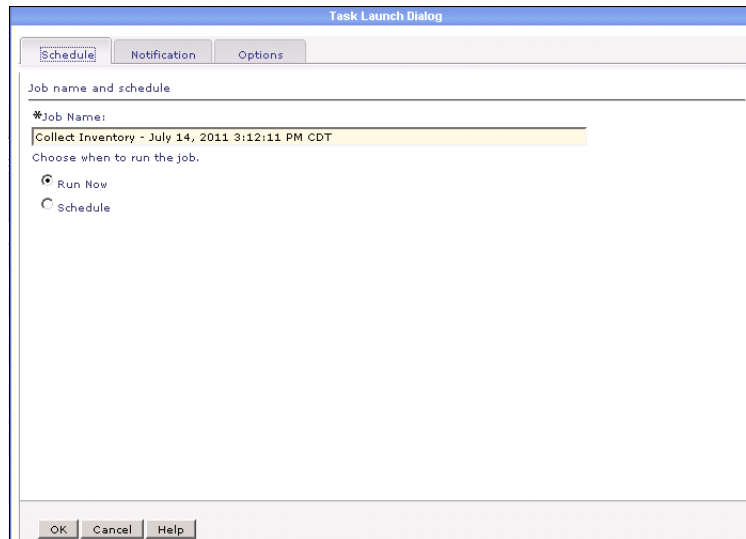


Figure 3-94 Task Launch Dialog

You can use Storage Control to monitor the health of storage subsystems, storage pools, storage volumes, and fabric switches.

Run the inventory against the Storage Control farm to see updated health status. Follow these steps to retrieve the health and alert status for the devices managed by Storage Control.

1. Select **Resource Explorer** → **All Systems**. Select the managed device for which you want to get the latest health and alert status.
2. Discover and then unlock the device.

Tip: At this point, you will receive TPC alerts from the farm, but you will not yet receive alerts specific to any devices under this farm.

3. Select the farm that covers the selected device and collect inventory against it.
4. For DS3000, DS4000, and DS5000 TPC devices only: Under the farm, right-click the DSnnnn device, and select Collect inventory.

Important: You must repeat this step each time you want to check for an event on these devices.

5. You will now receive TPC alerts for devices under the farm.

In this chapter we have covered management of the IBM Flex System V7000 Storage Node from CMM and FSM. For more information regarding CMM and FSM, see *Implementing Systems Management of IBM PureFlex System*, SG24-8060, and the IBM InfoCenter:

http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp?topic=%2Fcom.ibm.acc.pureflex.doc%2Fp7een_template_landing.html



IBM Flex System V7000 Storage Node initial configuration

This chapter provides the initial configuration steps of the IBM Flex System V7000 Storage Node. Before the actual installation and the initial configuration, proper planning is important. We also present an overview of the planning tasks and refer you to the appropriate documentation as needed to enable you to complete these tasks.

To help you get started and guide you to a successful installation and initial configuration of your IBM Flex System V7000 Storage Node system, the following topics are described:

- ▶ Planning and Installation requirements
- ▶ Initial setup with IBM Flex System Manager Node and IBM Flex System Chassis Management Module
- ▶ Setup Wizard for initial configuration
- ▶ System Management
- ▶ Service Assistant

4.1 Planning overview

We start with an overview of the planning tasks required for the proper installation of IBM Flex System V7000 Storage Node.

4.1.1 Hardware planning

Proper planning before the actual physical installation of the hardware is required. Here is an overview of the requirements to use to plan your installation:

- ▶ Each IBM Flex System V7000 Storage Node enclosure within the IBM Flex System Enterprise Chassis uses four bays (double wide and double high). Make sure that you have enough space in the chassis for the installation of IBM Flex System V7000 Storage Node control and internal expansion enclosures.
- ▶ IBM Flex System V7000 Storage Node control and expansion enclosures are attached in a single redundant SAS chain. Each IBM Flex System V7000 Control Enclosure will support up to nine expansion enclosures total, which constitutes an IBM Flex System V7000 Storage Node.
- ▶ IBM Flex System V7000 Control Enclosure supports the attachment for up to nine Expansion Enclosures using either IBM Flex System V7000 Expansion Enclosures (maximum 2), IBM Storwize V7000 Expansion Enclosures (maximum 9), or a combination of both.
- ▶ The IBM Flex System Enterprise Chassis contains 14 bays in total, so there can be at most three IBM Flex System V7000 Storage Node enclosures (a control enclosure and up to two IBM Flex System V7000 Expansion Enclosures). Any additional expansion enclosures would have to be IBM Storwize V7000 expansion enclosures attached externally to the same single SAS chain.

Tip: Installing a Flex System V7000 Enclosure within an Enterprise Chassis requires the removal of some of the shelves within the chassis, which would normally support the compute nodes.

Install the hardware as described in the *IBM Flex System V7000 Storage Node Installation Guide*, GC27-4203.

For information about planning the IBM Flex System V7000 Storage Node environment, see the following website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.4939.doc/site_product_page.html

For the latest maximum configuration support information, see the IBM System Storage Flex System V7000 Storage Node website:

http://www.ibm.com/support/entry/portal/overview/hardware/puresystems/pureflex_system/storage_node/flex_system_v7000

- ▶ There need to be two power outlets from redundant power sources in the rack for the two power cords included with each enclosure. From 2 to 20 outlets are required, depending on the number of enclosures to install. The power cord connectors are IEC320 C13/C14 standard connections.
- ▶ There are four optional configurations that can exist for host connections to the IBM Flex System V7000 Storage Node through network adapters. See Chapter 11, “SAN connections and configuration” on page 431 for detailed examples of how to configure the host connections.

- ▶ You should have a minimum of three IP addresses for initial system configuration. These IP addresses can be of either the IPv4 or IPv6 standards; one for management and two for service access to each of the control canisters.
- ▶ A minimum of one and up to four IP addresses are needed if iSCSI attached hosts access volumes from the IBM Flex System V7000 Storage Node. They are in addition to the three IP addresses needed for the management of the system.

4.1.2 SAN configuration planning

The SAN fabric is an area of the network that contains routers and switches. The IBM Flex System V7000 Storage Node has connections for up to eight 8 Gbps Fibre Channel connections per control enclosure. A SAN is configured into a number of zones. A device using the SAN can communicate only with devices that are included in the same zones that it is in. A Flex System V7000 Storage Node can have several distinct types of zones: a system zone, host zones, and disk zones. The intersystem zone is optional. Some operating systems cannot tolerate other operating systems in the same host zone, although you might have more than one host type in the SAN fabric. For example, you can have a SAN that contains one host that runs on an IBM AIX operating system and another host that runs on a Microsoft Windows operating system.

After the IBM Flex System V7000 Storage Node, hosts, and optional external storage systems are connected to the SAN fabrics, zoning needs to be implemented. See Chapter 11, “SAN connections and configuration” on page 431 for more information.

Tip: All communication between the canisters in the same control enclosure is performed either through the SAN or by PCI Express (PCIe) link. All configuration and service commands are sent to the system through an Ethernet network.

The recommended SAN configuration is composed of a minimum of two fabrics with all IBM Flex System V7000 Storage Node ports, host ports, and ports from external storage systems. They are to be virtualized as divided equally between the two fabrics for redundancy in the event that one of the fabrics goes offline (either planned or unplanned).

Degraded performance can occur when you have a SAN configuration in which any single component might fail and connectivity between the devices within the SAN is maintained. Splitting the SAN into two independent counterpart SANs achieves this normally. Because of the requirement for high availability, IBM Flex System V7000 Storage Node generally is installed into a redundant SAN, however, a non-redundant SAN is supported.

4.1.3 LAN configuration planning

The Enterprise Chassis provides the following services to the IBM Flex System V7000 Storage Node:

- ▶ A 1 Gbps management network
- ▶ Four host attachment networks:
 - 10 Gbps Ethernet
 - 8 Gbps or 16 Gbps Fibre Channel Service IP address

4.1.4 Management IP address considerations

The IBM Flex System V7000 Storage Node management IP address provides access to the system management interfaces, including the GUI and CLI. The management IP address is also used to access remote services such as authentication servers, NTP, SNMP, SMTP, and Syslog systems, if configured. The management IP address is configured during first-time setup of your IBM Flex System V7000 Storage Node. See 4.3, “IBM Flex System V7000 Storage Node Setup Wizard” on page 162.

The management IP address is associated with one of the control canisters in the IBM Flex System V7000 Control Enclosure and that control canister becomes the *configuration node*. Should this node go offline, either planned or unplanned, the management IP address does a failover to the surviving partner control canister’s Ethernet port and it becomes the *configuration node*.

Tip: IP addresses that are used by hosts to access the system over an Ethernet connection are different from management IP addresses.

Observe the following rules when configuring a management IP for your IBM Flex System V7000 Storage Node:

- ▶ Ethernet port 1 (1 Gbps) must be configured with an IPv4 or IPv6 system address.
- ▶ Ethernet port 2 (1 Gbps) can optionally be configured with a second system address.
- ▶ A maximum of one IPv4 address and one IPv6 address can be configured on each of Ethernet ports 1 and 2 for system addresses.
- ▶ To ensure system IP failover operations, Ethernet port 1 on all nodes must be connected to the same subnet. The system IP address can failover to any control node in the system.
- ▶ If Ethernet port 2 (1 Gbps) is configured with a system IP address, Ethernet port 2 on all nodes must also be connected to the same subnet. However, the subnet for Ethernet port 2 does not have to be the same as Ethernet port 1.
- ▶ Management addresses cannot be configured on ports 3 or 4 (10 Gbps), if present.

4.1.5 Service IP address considerations

The Service Assistant is a web application that helps troubleshoot and resolve problems in a Flex System V7000 Storage Node enclosure. Service Assistant is connected to on each IBM Flex System V7000 Control Enclosure node canister through the Service IP address. The node canister cannot be active as part of a system while it is in service state.

The Ethernet port 1 on each node canister is used for system management, and also used for service access when required. In normal operation, the Service IP addresses are not needed. However, in the event of a node canister problem, it might be necessary for service personnel to log on to the node to perform service actions.

Tip: The password is required to access the system through the Service IP address. The authentication of the superuser is always local; therefore, this user ID can be used even when a remote authentication server that is used for other users is not available.

To establish connectivity, the management node attempts to use Dynamic Host Configuration Protocol (DHCP) to acquire its initial IP address for the Chassis Management Module (CMM) Ethernet port. If DHCP is not installed or is enabled and fails, the management node uses the static IP address.

4.1.6 Management interface planning

IBM PureFlex Systems are pre-configured, pre-integrated infrastructure systems with different management options:

- ▶ Chassis Management Module (CMM)
- ▶ Flex System Manager (FSM)

For information and the actual implementation of the management modules, see *Implementing Systems Management of IBM PureFlex System*, SG24-8060.

Chassis Management Module (CMM)

The IBM Flex System Enterprise Chassis Management Module (CMM) is a hot-swap module that configures and manages all installed chassis components. The chassis comes with one Chassis Management Module in the CMM bays.

The IBM Flex System Chassis Management Module auto-detects and manages the hardware of the entire Flex chassis and provide centralized system monitoring and management.

Tip: The IBM Flex System Chassis Management Module provides control and information on the storage hardware, while the IBM Flex System Manager additionally provides configuration functionality.

For more details, see Chapter 3, “Systems management” on page 81.

The IBM Flex System Manager Node (FSM)

The Flex System Manager (FSM) is an advanced chassis management appliance that provides the ability to manage up to four enterprise chassis. It is designed with a specialized compute node, storage, networking, and provides physical and virtual management capabilities.

The FSM appliance is based on an Intel-based x86 compute node that comes with preloaded management software. The FSM software contains a fully operational management stack for managing multiple chassis from a single focal-point for support of all assets and resources in IBM PureFlex Systems handling the complete range of management tasks. The following list of the tasks can be performed from the FSM:

- ▶ Service requests
- ▶ Incident handling and problem resolution
- ▶ Handling specialized management tasks in connection with the following functions:
 - Security
 - Performance
 - Resource allocation
 - Application version control

4.2 Initial setup for IBM Flex System V7000 Storage Node

When a new IBM Flex System V7000 Storage Node control enclosure is installed in the Flex System chassis for the first time, you must create and configure it into a *clustered system*. The following procedure guides you through the necessary steps to set up your IBM Flex System V7000 Storage Node system for the first time in this configuration.

Choosing a method

There are two methods that can be used for the initial setup of the IBM Flex System V7000 Storage Node. The method used depends upon the configuration of the IBM Flex System:

- ▶ If the IBM Flex System has an installed and configured IBM Flex System Manager, then it should be used to set up the IBM Flex System V7000 Storage Node.
- ▶ If the IBM Flex System does not have an IBM Flex System Manager (FSM) installed and configured, use the Chassis Management Module to set up the IBM Flex System V7000 Storage Node.

Before starting

Observe these preliminary considerations:

- ▶ You should have installed the IBM Flex System V7000 Storage Node enclosures in the rack, connected all cables, powered the system on, and checked the LED status for the system. When you have completed these steps, you can continue with the configuration of the IBM Flex System V7000 Storage Node.

- ▶ To complete this procedure, use a supported web browser. Verify that you are using a supported web browser from the following website:

http://www.ibm.com/support/entry/portal/overview/hardware/puresystems/pureflex_system/storage_node/flex_system_v7000

- ▶ Decide whether you are going to use an DHCP server to assign the node service IP address or will assign static IP addresses.

If using an DHCP server, make sure the IBM Flex System V7000 Storage Node has access to the DHCP server and there are at least two DHCP addresses available on the subnet for their use. They are the addresses that will be used in the initial installation.

If assigning static service IP addresses, make sure that you also know the management IP address that you want to assign to the IBM Flex System V7000 Storage Node. It is the address that will be used to access the IBM Flex System V7000 Storage Node cluster for management after installation is complete.

- ▶ If you are using the Chassis Management Module (CMM) to set up your system, ensure that the chassis the IBM Flex System V7000 Storage Node is installed in is being managed by the CMM.

If you are planning to use the CMM to setup your system, verify that you are using a CMM that supports the IBM Flex System V7000 Storage Nodes from the following website:

http://www.ibm.com/support/entry/portal/overview/hardware/puresystems/pureflex_system/storage_node/flex_system_v7000

- ▶ If you are using the Flex System Manager (FSM) to setup your system, you must have an installed and pre-configured FSM. Verify that you are using a supported FSM for IBM Flex System V7000 Storage Nodes from the following website:

http://www.ibm.com/support/entry/portal/overview/hardware/puresystems/pureflex_system/storage_node/flex_system_v7000

Tip: A USB key is not required for first-time setup or normal operations. However, it might be required for resetting the password, the service IP, and to check the status of an inaccessible canister, if directed by IBM Service and Support.

4.2.1 Using FSM for initial setup

Now that you have installed the IBM Flex System V7000 Storage Node in the chassis with supported level of code, you can create and configure a clustered system using the FSM. The following procedure guides you through the necessary steps when using the FSM web user interface:

1. Open a web browser and point it to the IP address of the FSM and log in. The following menu panel (Figure 4-1) displays, giving you a number of selections.

Select **Launch IBM FSM Explorer** from the menu list.

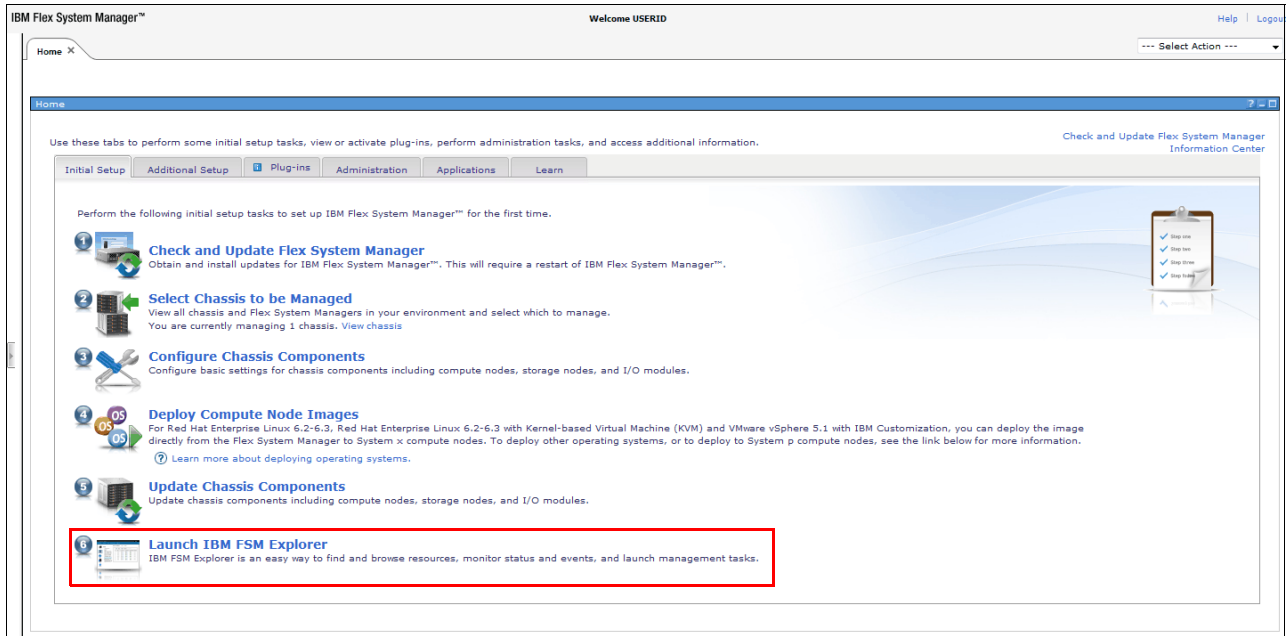


Figure 4-1 Launch IBM FSM Explorer

Notice that a new browser tab is opened, which allows you to select the applicable enclosure from the Chassis Map as shown in Figure 4-2.

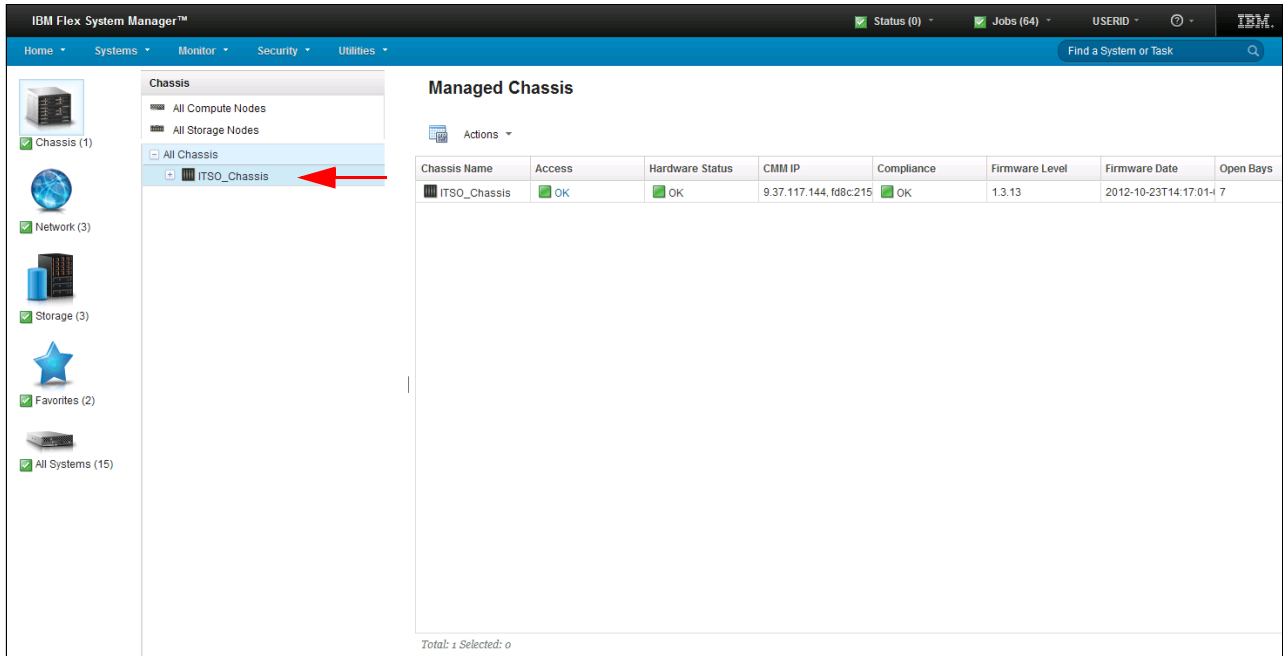


Figure 4-2 Select and launch the chassis in the Chassis Manager

2. In the Chassis Manager, select the applicable chassis which will launch the Chassis Map for that chassis, as shown in Figure 4-3.




Figure 4-3 IBM Flex System Manager - Hardware Map

3. From the Chassis Manager page in the management software web interface, click the chassis name that contains your applicable storage node as shown in Figure 4-4 (A) and continue to Step 4. If the chassis is not displayed, add the chassis by performing the following steps:
 - a. Select **General Actions** → **Add or Remove a chassis**, as shown in Figure 4-4 (B).
 - b. Select the chassis that you want to manage.
 - c. Click **Manage**. The Management Domain page displays.
 - d. Click **Manage**.
 - e. Enter the user ID and password for the CMM of the chassis that you want to manage.
 - f. Wait while the status is processed. After the status is processed, the status changes to Success, and the message changes to Managed.
 - g. Click **Done**.




Figure 4-4 Add a chassis into Chassis Manager

4. Navigate to General Actions and click **Launch IBM Flex System V7000 Storage Node Manage (V7000)** as shown in Figure 4-5 to start the Initial Setup wizard. If you do not see this option, examine the Service IP setting for the storage node. For details, see 4.5.1, “Changing the Service IP address” on page 171.

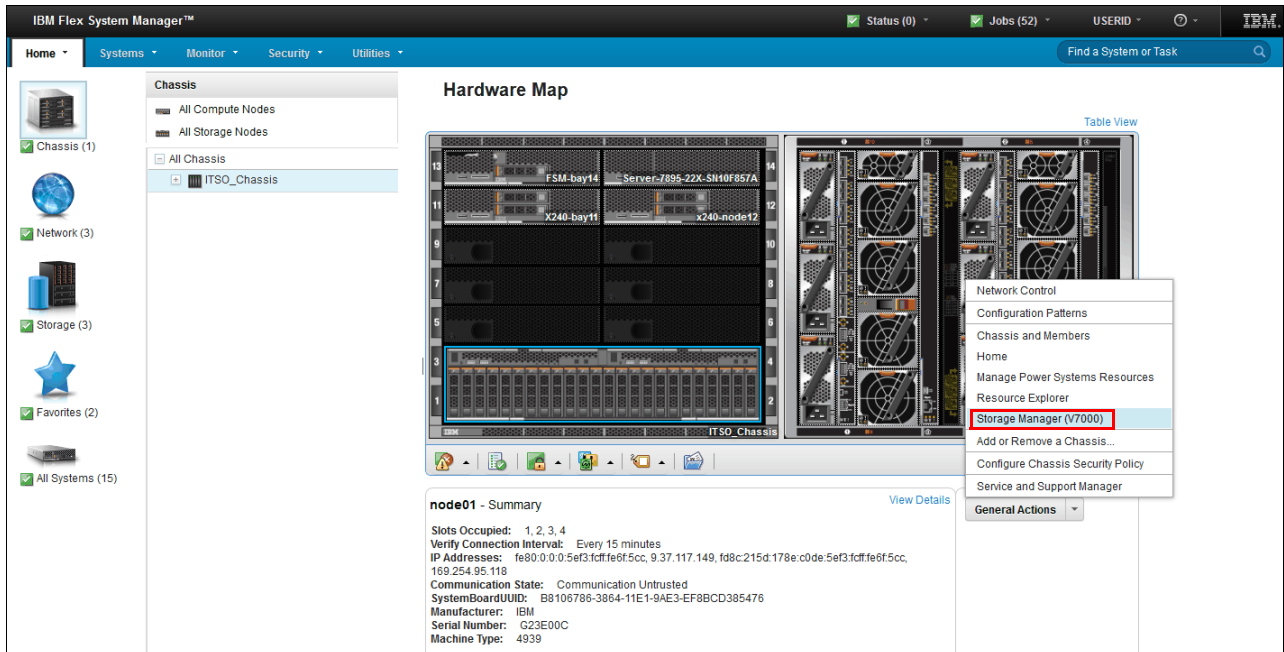


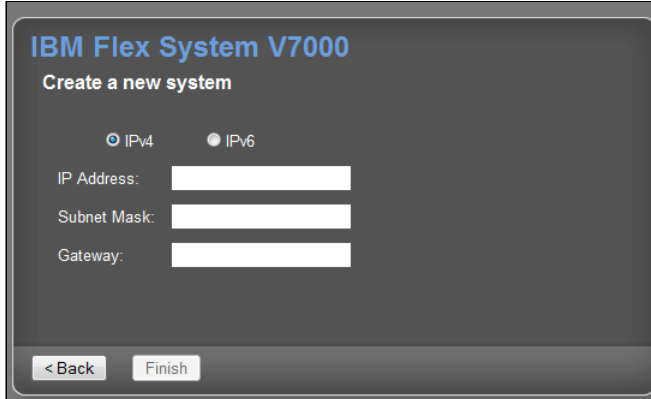
Figure 4-5 Launch Storage Manager (V7000)

5. The next window is a welcome window from the IBM Flex System V7000 Storage Node interface, asking to either create a new system (cluster) or add to an existing system, as shown in Figure 4-6.



Figure 4-6 IBM Flex System V7000 Storage Node first time setup welcome window

6. If you choose to create a new system, it will ask for network information, as shown in Figure 4-7. Select whether you are using an IPv4 or IPv6 management IP address and type in the address (you can use either DHCP or the static address that was assigned). The subnet mask and gateway will already have defaults listed, which you can edit.



The screenshot shows a web-based configuration window titled "IBM Flex System V7000" with the subtitle "Create a new system". It features two radio buttons: "IPv4" (which is selected) and "IPv6". Below these are three text input fields labeled "IP Address:", "Subnet Mask:", and "Gateway:". At the bottom of the window, there are two buttons: "< Back" and "Finish".

Figure 4-7 Create new Storage Cluster

7. Click **Finish** to set the management IP address for the system. System initialization begins and might take several minutes to complete.

When system initialization is complete, System Setup is launched automatically. The setup wizard will take you through the steps to configure basic system settings such as time and date, system name, and hardware detection and verification.

For more details on the rest of the setup wizard steps, see 4.3, “IBM Flex System V7000 Storage Node Setup Wizard” on page 162.

4.2.2 Using CMM for initial setup

When the first control enclosure for a new IBM Flex System V7000 is installed in the Flex System chassis, you must create and configure a system. This procedure is used to guide you through the necessary steps using the CMM for the first time.

Before you begin

Here are some considerations before you begin the setup:

- ▶ You need to be sure that your IBM Flex System V7000 Storage Node has access to the DHCP server and there are at least two DHCP addresses available on your subnet. They are the addresses that will be used in the initial installation. If DHCP addresses are unavailable, the node will use a default service IP address based on the slot location. It can be changed to a static address using the CMM Component IP Configuration page.
- ▶ When a new control enclosure is installed in the IBM Flex System Enterprise Chassis, the CMM automatically recognizes the new hardware. Open the CMM web user interface and navigate to the Chassis Map. In the Chassis Map displayed by the CMM, you will see the new Flex System V7000 Storage Node.

Initial setup

Follow this procedure for the initial setup from the CMM interface:

1. Log in to the CMM and navigate to the Chassis Map. In the Chassis Map displayed by the CMM, you will see all the IBM Flex System V7000 Storage Node enclosures installed in the chassis.

Tip: If you have installed more than one control enclosure in your IBM Flex System V7000 Storage Node system, choose a control enclosure to be used to initialize the system. It does not matter which control enclosure you choose. You cannot initialize the system from an expansion enclosure.

2. Select the **Actions** menu and click **Launch Storage Node Console**. The Action menu can be found by right-clicking the left canister in the graphical representation of the storage node, as shown in Figure 4-8.

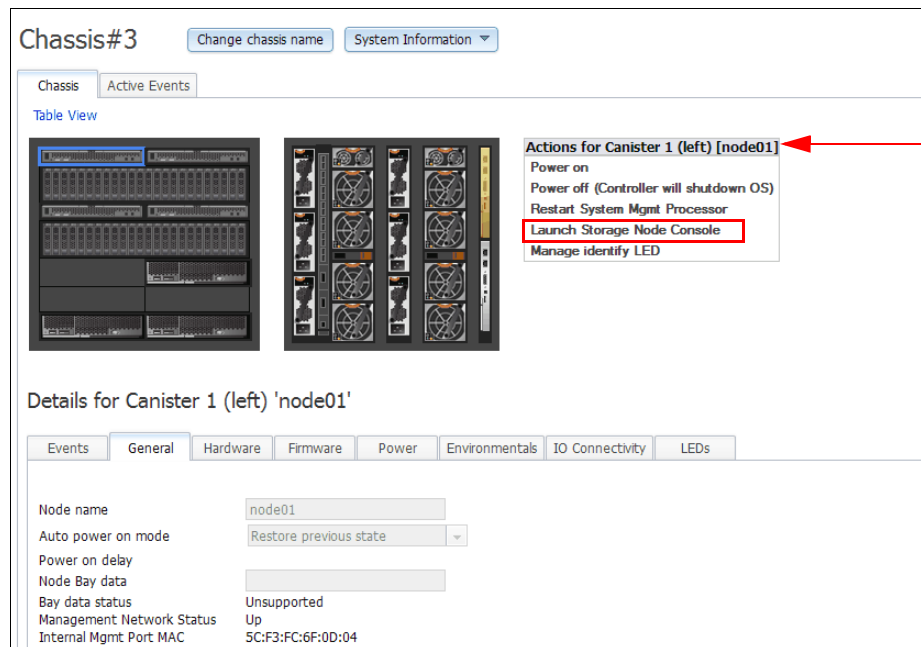


Figure 4-8 CMM - Select Launch Storage Node Console

3. A browser window displays the IBM Flex System V7000 Storage Node system startup panel. The Service IP address displayed will be one that has been allocated by your DHCP server as shown in Figure 4-9.

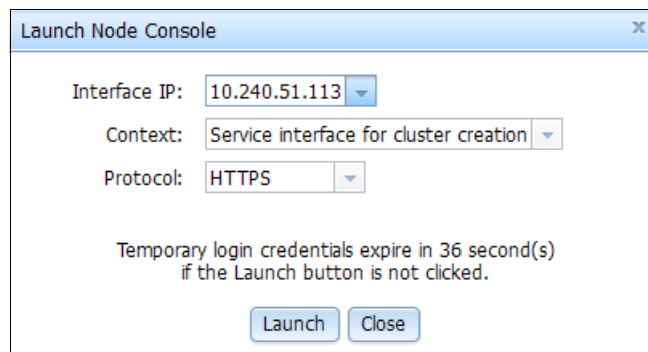


Figure 4-9 CMM - Launch Node Console

- If a default Service IP address in the range between 192.168.70.131 and 192.168.70.144 is displayed, the canister has not obtained a DHCP address. This default service IP address can be changed to a static address using the CMM Component IP Configuration page. See the topic “Changing the service IP address of a node canister” in the *IBM Flex System V7000 Storage Node Troubleshooting, Recovery, and Maintenance Guide* PDF on the CD provided with your system, or in the IBM Flex System V7000 Storage Node Information Center.
- If you have problems connecting to the service address, see the Cannot connect to the service assistant topic in the *IBM Flex System V7000 Storage Node Troubleshooting, Recovery, and Maintenance Guide* PDF on the CD or in the IBM Flex System V7000 Storage Node Information Center.

After clicking **Launch** in Figure 4-9, you have an option to Create a new system or Add to an existing system as shown in Figure 4-10. We select **Create a new system**, then click **Next**.



Figure 4-10 IBM Flex System V7000 Storage Node first time setup welcome window

4. Select whether you are using an IPv4 or IPv6 management IP address and type in the address (you can use DHCP or statically assigned one). The subnet mask and gateway will already have defaults listed, which you can edit, as shown in Figure 4-11.




Figure 4-11 Create new Storage Cluster

5. Click **Finish** in Figure 4-11 to set the management IP address for the system. System initialization begins and might take several minutes to complete.

When system initialization is complete, System Set Up is launched automatically. The setup wizard will take you through the steps to configure basic system settings such as time and date, system name, and hardware detection and verification.

For more details on the rest of the setup wizard steps, see 4.3, “IBM Flex System V7000 Storage Node Setup Wizard” on page 162.

4.3 IBM Flex System V7000 Storage Node Setup Wizard

After the initial configuration using the FSM or CMM, we can continue with the setup wizard for the rest of the configuration of IBM Flex System V7000 Storage Node.

After clicking **Finish** in Figure 4-11 on page 161, on the next panel, enter the system name and a new superuser password. At this point, you have the option to use the setup wizard to configure more settings such as notifications and storage configuration. If you are not ready to complete more configuration steps now, use the configuration tasks in the management GUI to complete the configuration at a later time.

When you complete the Initial Setup wizard, return to the FSM or CMM console.

Tip: It is recommended that you upgrade to the most current level of software after installing the Flex System V7000 Storage Node. The management GUI can be used to apply software updates. See the IBM Flex System V7000 Storage Node support website for the latest information about software upgrades.

1. After the initial configuration described in 4.2, “Initial setup for IBM Flex System V7000 Storage Node” on page 153 is complete, the IBM Flex System V7000 Storage Node Welcome window opens (Figure 4-12).



Figure 4-12 IBM Flex System V7000 Storage Node Welcome window

Tip: During the initial setup of the Flex System V7000, the installation wizard asks for various information that you should have available during the installation process. If you do not have this information or choose not to configure some of the items at this time, you can configure them later through the GUI.

2. You have to read and accept the license agreement to be able to continue as shown in Figure 4-13. Click **Next** after accepting the licence agreement.

License Agreement (Step 1 of 7)

Read the license agreement carefully.

License IBM Notices Java Notices Non-IBM Licenses Additional Licenses and Notices

International Program License Agreement

Part 1 - General Terms

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

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

* PROMPTLY RETURN THE UNUSED MEDIA, DOCUMENTATION, AND PROOF OF ENTITLEMENT TO THE PARTY FROM WHOM IT WAS OBTAINED FOR A REFUND OF THE AMOUNT PAID. IF THE PROGRAM WAS DOWNLOADED, DESTROY ALL COPIES OF THE PROGRAM.

1. Definitions

"Authorized Use" - the specified level at which Licensee is authorized to execute or run the Program. That level may be measured by number of users, millions of service units ("MSUs"), Processor Value Units ("PVUs"), or other level of use specified by IBM.

"IBM" - International Business Machines Corporation or one of its subsidiaries.

I agree with the terms in the license agreement.

I do not agree with the terms in the license agreement.

Next >

Figure 4-13 Setup wizard - License Agreement

3. Set up the System Name and the Superuser Password and click **Next** as shown in Figure 4-14.

System Name and Superuser Password (Step 2 of 7)

System Name

* Name: FlexSystem_V7000

Superuser Password

New Superuser Password

Verify New Superuser Password

Next >

Figure 4-14 Setup wizard - Set system name and superuser password

4. Set up the system date and time and click **Next** as shown in Figure 4-15.

Date and Time (Step 3 of 7)

Current Date and Time

23/07/2012 11.06.37

Time Zone

(GMT-5:00) US Eastern Time

Use Browser Settings

Next >

Figure 4-15 Setup wizard - Set date and time

- Optionally, you can type in System licenses and click **Next** as shown in Figure 4-16. The System Licenses include External Virtualization Limit, Remote-Copy Limit, and Real-time Compression Limit. The virtualization license for all directly attached expansion enclosures are already included in the System License and should not be added here.

System License (Step 4 of 7)

The enclosure license already includes virtualization of internal Serial Attached SCSI (SAS) drives on your system. You can use this panel to set any additional options. If you are sharing the total authorized capacities across multiple clusters, enter only the capacities you wish to use on this system. The sum of the capacities across all systems must not exceed your authorized capacities.

Set License Options

External Virtualization Limit
 enclosures

Remote-Copy Limit
 enclosures

Real-time Compression Limit
 enclosures

Next >

Figure 4-16 System license

- Configure support notifications as shown in Figure 4-17.

Configure Support Notifications (Step 5 of 7)

To ensure your system continues to run smoothly you can enable email event notifications.

Email event notifications send messages about error, warning, or informational events and inventory reports to an email address of local or remote support personnel. Ensure that all the information is valid, otherwise email notification will be disabled.

Configure Email Event Notifications

[Watch e-Learning: Configure Support Notifications](#)

Next >

Figure 4-17 Configure support notifications

- Define company contact information as shown in Figure 4-18.

Configure Support Notifications Step 1 of 4

Define Company Contact

Support personnel can contact this person to assist with problem resolution. Ensure that all contact information is valid.

Email Contact

* **Contact Name** * **Email Reply Address**

* **Machine Location** * **Telephone (Primary)** **Telephone (Alternate)**

* Required

Next > **Cancel**

Figure 4-18 Define company contact

8. Verify that all hardware has been detected by the system correctly as shown in Figure 4-19.

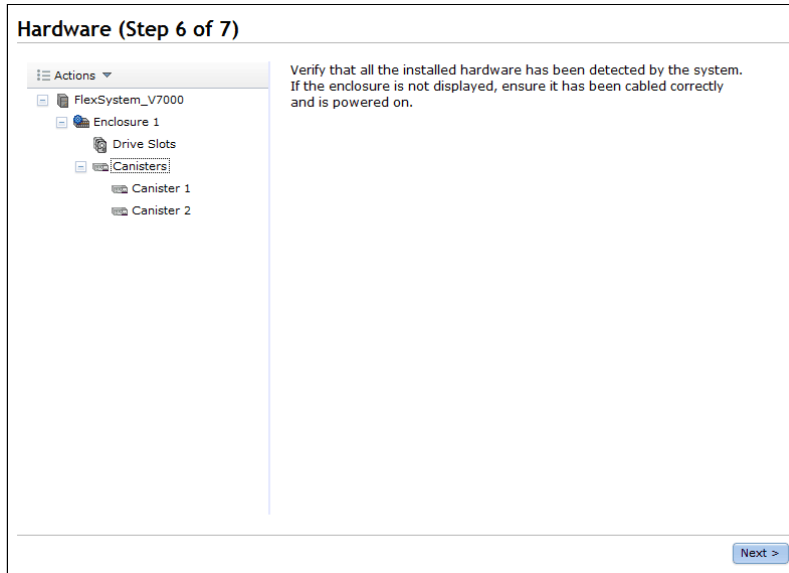


Figure 4-19 Verify hardware

9. Select **Yes** to automatically configure internal storage now if preferred (Figure 4-20).

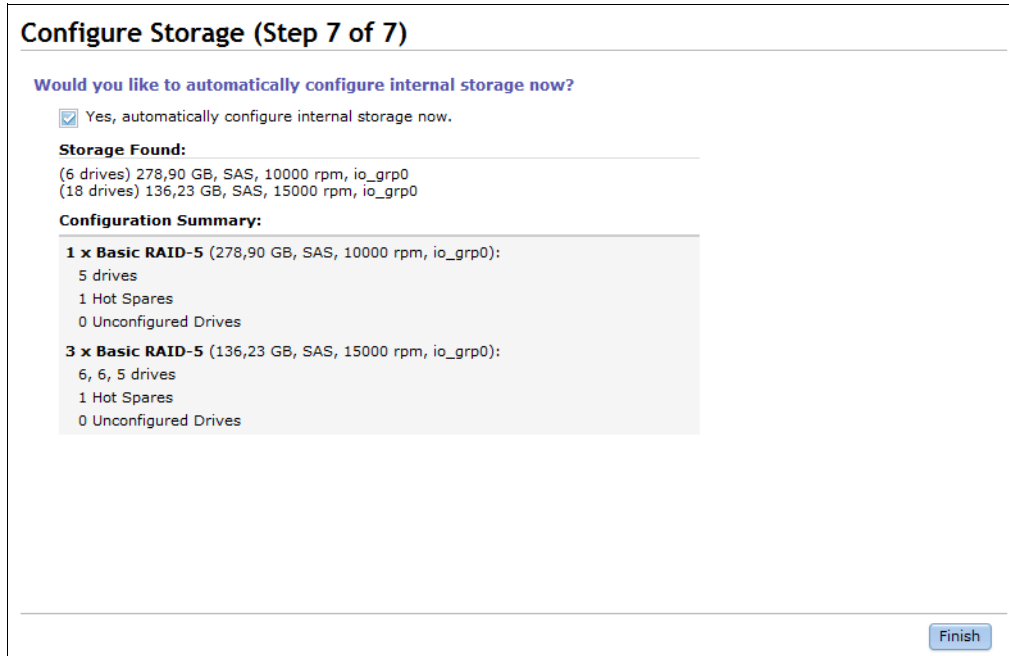


Figure 4-20 Select “Yes” to automatically configure internal storage now

10. Click **Finish** to complete the Setup wizard task and log in to IBM Flex System V7000 Storage Node as shown in Figure 4-21. You log in as superuser with your newly defined password. If you have not changed the password, the default is “passw0rd”.



Figure 4-21 Setup wizard task complete

11. After successful login, the IBM Flex System V7000 Storage Node Home Overview window looks as shown in Figure 4-22.

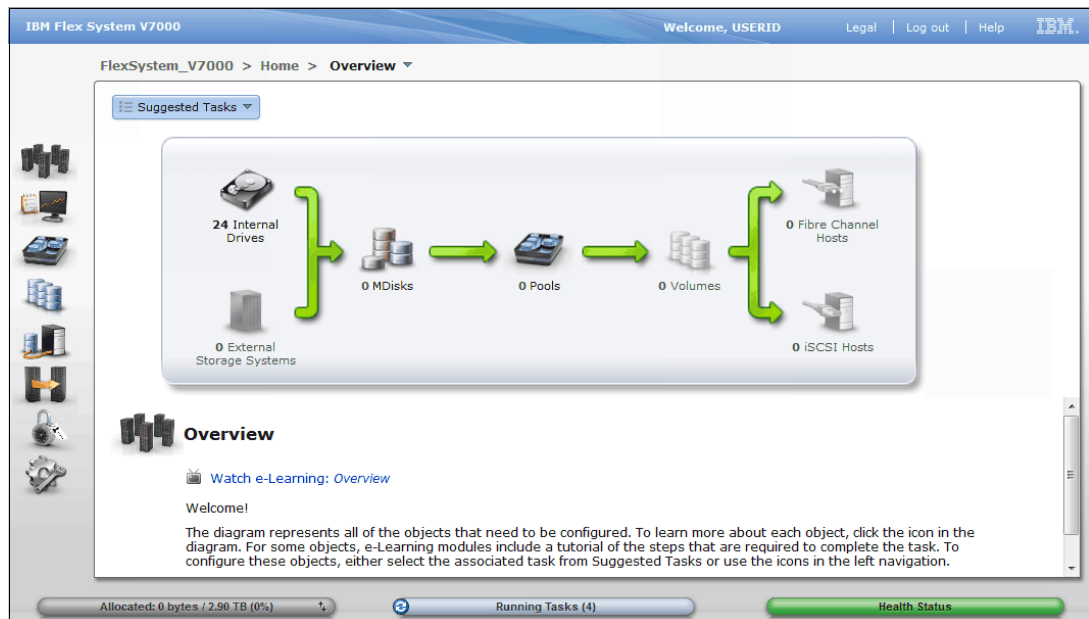


Figure 4-22 IBM Flex System V7000 Storage Node Home Overview window

12. IBM Flex System V7000 Storage Node initial configuration is complete and the cluster is up and running (Figure 4-23).

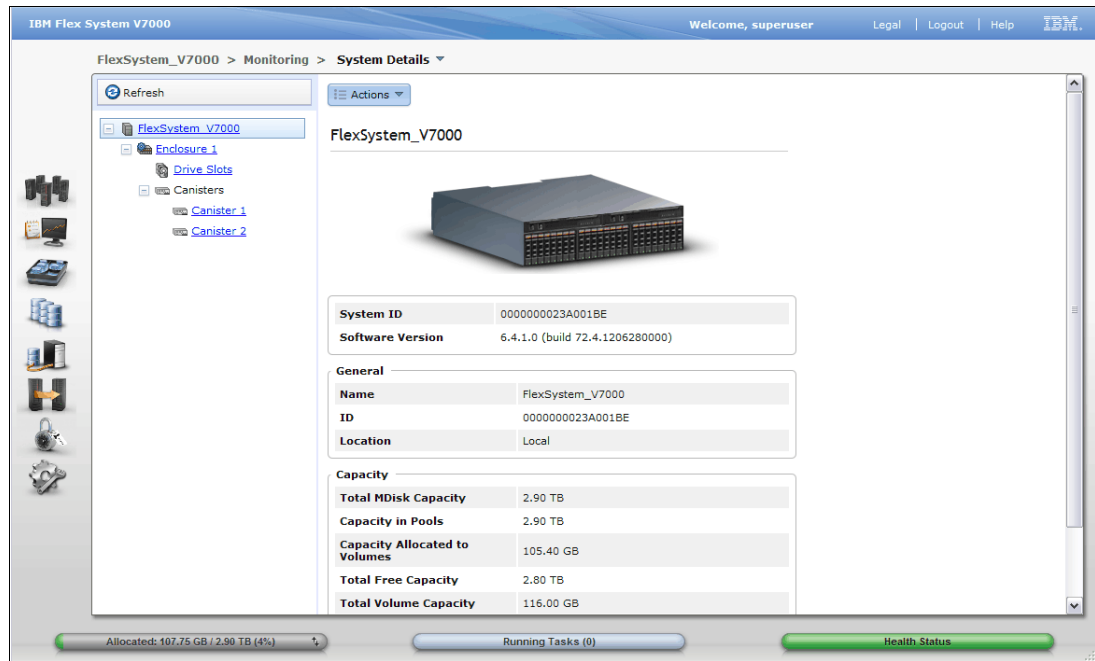


Figure 4-23 System details view in management GUI

13. You can continue to configure additional functions and features for your environment in order to meet your implementation requirements.

4.4 System management

System management allows for troubleshooting and management tasks using the management graphical user interface (GUI). It includes checking the status of the storage server components, updating the firmware, monitoring events, and managing the storage server. It offers advanced functions, such as FlashCopy, Volume Mirroring, and Remote Mirroring.

A command-line interface (CLI) for the IBM Flex System V7000 Storage Node is also available.

IBM Flex System V7000 Storage Node is managed through FSM or CMM using the management GUI or a command-line interface (CLI) over an Ethernet connection. For more information, see Chapter 3, "Systems management" on page 81.

Tip: Management IP addresses that are assigned to a Flex System V7000 Storage Node are different from iSCSI IP addresses and are used for different purposes. If iSCSI is used, iSCSI addresses are assigned to node ports. On the configuration node, a port has multiple IP addresses active at the same time.

4.4.1 Graphical User Interface (GUI)

The IBM Flex System V7000 Storage Node graphical user interface (GUI) is designed to simplify storage management and provide a fast and more efficient management tool. Use a supported web browser for GUI access to configure, manage, and troubleshoot the IBM Flex System V7000 Storage Node. It is also used primarily to configure RAID arrays and logical drives, assign logical drives to hosts, replace and rebuild failed disk drives, and expand the logical drives.

JavaScript: You might need to enable JavaScript in your browser. Additionally, if you are using Firefox, under Advanced JavaScript Settings, you need to click **Disable** or replace context menus and allow cookies.

After the first-time setup procedure is complete, open your web browser on a workstation and point it to your IBM Flex System V7000 Storage Node management IP address (Figure 4-24) and complete the steps guided by the Initial Setup wizard. The default login is superuser and the password is passw0rd.

Tip: If a management IP has been configured for your IBM Flex System V7000 Storage Node, the system name that you specified will appear in the login window, as shown in Figure 4-24. If a longer system name is used, you might need to hover over the system name to view the full system name, due to limited display space.



Figure 4-24 IBM Flex System V7000 Storage Node GUI logon panel

Figure 4-25 shows an overview of the IBM Flex System V7000 Storage Node GUI.



Figure 4-25 IBM Flex System V7000 Storage Node management GUI Welcome page

For more details on how to use IBM Flex System V7000 Storage Node management GUI, see Chapter 5, “IBM Flex System V7000 Storage Node GUI interface” on page 175.

4.4.2 Launching IBM Flex System V7000 Storage Node GUI from CMM

When a new control enclosure is detected in the IBM Flex System Enterprise Chassis, the Chassis Management Module (CMM) recognizes the new hardware and starts the initialization process. The IBM Flex System V7000 Storage Node Management GUI can be launched and accessed by CMM after the management IP been assigned to your new Flex System V7000 storage node.

To start, log in to the CMM GUI and navigate to the Chassis Map. Right-click the canister that you want to launch and select **Launch Storage Node Controller Console** to access the Storage Node management GUI. See Figure 4-26.

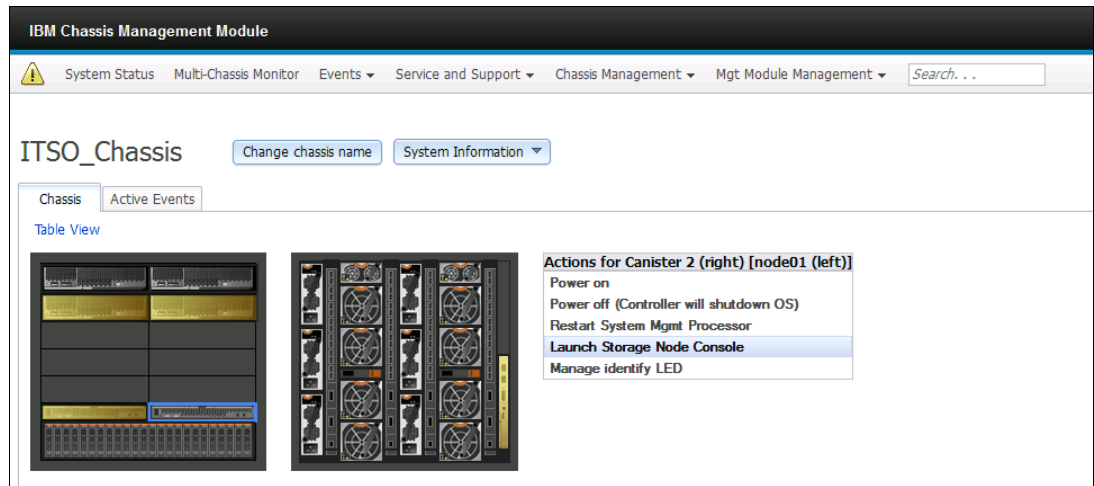


Figure 4-26 To launch storage node controller console by Chassis Map

To launch the IBM Flex System V7000 Storage Node management GUI, select the Interface IP address (Figure 4-27) and click **Launch** to start the Storage Node Controller Console.

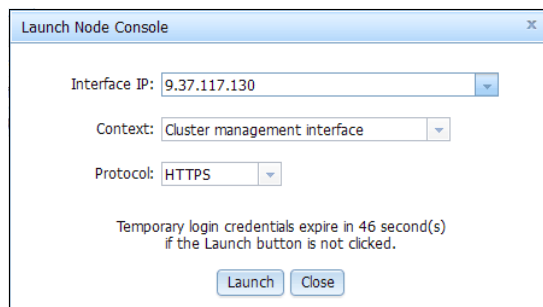


Figure 4-27 Launch to start the Storage Node Controller Console

4.5 Service Assistant

The primary use of the Service Assistant is when a node canister in the IBM Flex System V7000 Storage Node enclosure is in service state. The node canister might be in service state because it has a hardware issue, has corrupted data, or has lost its configuration data. The node canister cannot be active as part of a system while it is in service state.

Attention: Perform service actions on node canisters only when directed to do so by the fix procedures. If used inappropriately, the service actions that are available through the Service Assistant can cause loss of access to data or even data loss.

The storage system management GUI operates only when there is an online system. Use the Service Assistant if you are unable to create a system or if all node canisters in a system are in service state. Use the Service Assistant in the following situations:

- ▶ When you cannot access the system from the management GUI and you cannot access the Flex System V7000 Storage Node to run the recommended actions
- ▶ When the recommended action directs you to use the Service Assistant.

Tip: The Service Assistant does not provide any facilities to help you service expansion enclosures. Always service the expansion enclosures by using the management GUI.

You can also perform the following service-related actions using Service Assistant:

- ▶ Collect logs to create and download a package of files to send to support personnel.
- ▶ Remove the data for the system from a node.
- ▶ Recover a system if it fails.
- ▶ Install a software package from the support site or rescue the software from another node.
- ▶ Upgrade software on node canisters manually versus performing a standard upgrade procedure.
- ▶ Configure a control enclosure chassis after replacement.
- ▶ Change the Service IP address that is assigned to Ethernet port 1 for the current node canister.
- ▶ Install a temporary SSH key if a key is not installed and CLI access is required.
- ▶ Restart the services used by the system.

4.5.1 Changing the Service IP address

There are multiple methods that you can use to change the Service IP address of a storage node canister and which method to use depends on the status of the system and the other node canisters in the system. Follow the methods in the order shown until you are successful in setting the IP address to the required value.

Tip: You can set an IPv4 address, an IPv6 address, or both, as the service address of a node. Enter the required address correctly. If you set the address to 0.0.0.0 or 0000:0000:0000:0000:0000:0000, you disable the access to the port on that protocol.

Changing the Service IP address by management GUI

Use the control enclosure management GUI when the system is operating and the system is able to connect to the node canister with the Service IP address that you want to change, as shown in Figure 4-28.

1. Select **Settings** → **Network** from the navigation.
2. Be sure to select the correct node to configure.
3. Select **Specify IP address** and change it to the new IP.
4. Click **Save** to complete the panel.




Figure 4-28 Change Service IP by management GUI procedure

Changing the Service IP address by CMM

If the CMM in the chassis is operating and reachable, you can use the CMM to change the Service IP address of a node canister, as shown in Figure 4-29.

Tip: If the canister has failed to obtain a DHCP address, a default Service IP address in the range between 192.168.70.131 and 192.168.70.144 is displayed.

1. Open the CMM and navigate to the **Chassis Management** → **Component IP Configuration** page
2. In the list of Storage Nodes, click the name of the node canister.
3. In the IP address Configuration panel, click the **IPv4** tab.
4. Select **Node_Service_If** as the Network Interface and select the Configuration Method you want to use. If the selected Configuration Method is to Use Static IP address (as shown in Figure 4-29), type the Static IP address, Subnet Mask, and Gateway address for the New Static IP Configuration.
5. Click **Apply** and then **Close**.




Figure 4-29 Change Service IP by CMM procedure

Changing the Service IP address by CLI

These addresses are not set during the installation of a Flex System V7000 Storage Node system, but you can set these IP addresses later by using the `chserviceip` CLI command.

4.6 Command-Line interface (CLI)

The Flex System V7000 Storage Node command-line interface (CLI) is a collection of commands that you can use to manage a Flex System V7000 Storage Node.

You can use the CLI to perform the following functions:

- ▶ Setup of the clustered system, its nodes, and the I/O groups
- ▶ Analyzing error logs (event logs)
- ▶ Setup and maintenance of managed disks (MDisk) and storage pools
- ▶ Setup and maintenance of client public SSH keys on the clustered system
- ▶ Setup and maintenance of volumes
- ▶ Setup of logical host objects
- ▶ Mapping volumes to hosts
- ▶ Navigating from managed hosts to volumes and MDisks, and reverse direction up chain.
- ▶ Setting up and starting Copy Services:
 - FlashCopy and FlashCopy consistency groups
 - Synchronous Metro Mirror and Metro Mirror consistency groups
 - Asynchronous Global Mirror and Global Mirror consistency groups

The CLI commands use the Secure Shell (SSH) connection between the SSH client software on the host system and the SSH server on the Flex System V7000 Storage Node. Install and set up SSH client software on each system that you plan to use to access the CLI. Authenticate to the system using a password. If you require command line access without entering a password, use an SSH public key. Then store the SSH public key for each SSH client on the Flex System V7000 Storage Node.

Public Keys: After the first SSH public key is stored, you can add additional SSH public keys using either the management GUI or the CLI.

You can connect to the Flex System V7000 Storage Node CLI using PuTTY (Figure 4-30).

```
login as: superuser
superuser@9.37.117.130's password:
Last login: Mon Jul 23 21:08:45 2012 from 9.80.11.49
IBM_4939:FlexSystem_V7000:superuser>sainfo lsservicenodes
panel_name cluster_id      cluster_name  node_id node_name relation node_status error_data
01-1        0000000023A001BE FlexSystem_V7000 1      node1    local    Active
01-2        0000000023A001BE FlexSystem_V7000 2      node2    partner Active
IBM_4939:FlexSystem_V7000:superuser>
```

Figure 4-30 Connect to the CLI using PuTTY

A detailed Command Line Interface setup procedure for the IBM Flex System V7000 Storage Node is available in Appendix A, “CLI setup and configuration” on page 571. It can also be found at the following address:

http://pic.dhe.ibm.com/infocenter/storwize/ic/index.jsp?topic=/com.ibm.storwize.v7000.doc/svc_clicommandscontainer_229g0r.html

4.7 Recording system access information

It is important that anyone who has responsibility for managing the system knows how to connect to and log on to the system. Pay attention to those times when the normal system administrators are not available because of vacation or illness. Record the following information and ensure that authorized people know how to access the system information as shown in Table 4-1.

Table 4-1 Record log on information

Item	Value	Notes
Management IP address for the GUI and CLI		
Management user ID (the default is admin)		
Management user ID password (the default is admin)		
Control enclosure management IP address		
Control enclosure Service IP address: node canister 1		
Control enclosure Service IP address: node canister 2		
Control enclosure superuser password (default is passw0rd)		