System p and eServer pSeries

Clustering systems using InfiniBand hardware

System p and eServer pSeries

Clustering systems using InfiniBand hardware

Note

Before using this information and the product it supports, read the information in Appendix A, "Notices," on page 231 and the manual *IBM Systems Safety Information*, G229-9054.

Third Edition (September 2007)

© Copyright International Business Machines Corporation 2004, 2007. All rights reserved.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Clustering systems using InfiniBand hardware 1
Information resources 2
InfiniBand switch reference information 3

Planning for InfiniBand networks 5
Planning overview 5
Required levels of support, firmware, and devices . . 7
Planning InfiniBand network cabling and configuration 8
Management subsystem planning 10
Installation flow of clusters using the IBM Network Manager 11
Planning for a high-performance computing message-passing interface configuration 18
Planning octopus cables in static 12x cabling . . . 19
Planning 4x connections for octopus cables 21
Planning Aids 22
Planning checklist 22
Planning worksheets 23

Installing a cluster that has an InfiniBand network 39
Installing or replacing a GX Host Channel Adapter 64
Installing a GX Dual-Port Host Channel Adapter in a model 52A, 550, 550Q, 560Q, or 570 66
Installing a GX Dual-Port Host Channel adapter in a model 575, 590, or 595 66
Verifying the installed InfiniBand network (fabric) in AIX or Linux 67
Verifying the GX HCA connectivity in AIX . . . 67
Verifying the GX HCA to InfiniBand fabric connectivity in Linux 67
Verifying static-12x mode connectivity 67

Managing InfiniBand networks with IBM Network Manager 71
Enabling and disabling the IBM Network Manager 71
Viewing switch topology information in an InfiniBand network 71
Viewing server topology information in an InfiniBand network 72
Viewing logical topology information in an InfiniBand network 73
Viewing IBM Network Manager properties 74
Viewing the IBM Network Manager event log . . . 75
Updating switch software 75
Updating switch software from 2.3.0 or earlier to 2.5.0 or later 76

InfiniBand network problem determination 77

InfiniBand fabric maintenance strategy 81
Setting up a clustered environment to connect to service and support 81
Task 1. Before you begin 81
Task 2. Determine your connectivity method . . 81
Task 3. Prerequisites 82
Task 4. Ensure that your physical networking is set up correctly 82
Task 5. Obtain or verify an IBM ID 84
Task 6. Verify the HMC service settings using the Guided Setup wizard 84
Task 7. Set up and configure your logical partitions 86
Task 8. Install the operating systems on your server or logical partitions 86
Task 9. Configure your TCP/IP network. 86
Task 10. Activate TCP/IP on your server or logical partitions. 86
Task 11. Configure AIX or Linux for connectivity 86
Task 12. Use the Service Agent (SA) Basic User Interface 88
Task 13. Register the ID with Electronic Service Agent for AIX or Linux 89
Task 14. Configure the service processor. 89
Task 15. Test the connection to service and support. 90
Task 16. View the server information that was reported to IBM 91
Getting fixes for a clustered environment 91

Service information and procedures for networks managed by the IBM Network Manager 93
Reference codes for GX HCA-based InfiniBand networks 94

Miscellaneous service procedures for networks managed by the IBM Network Manager 99
Recovering from an HCA preventing a logical partition from activating. 100
Diagnosing an InfiniBand switch that will not boot 101
InfiniBand problem isolation procedures 101
IBNNMD. 102
IBNNURM 102
IBNSAUT 105
IBNSBAT. 105
IBNSCFG. 106
IBNSCRD 107
IBNSDBS. 107
IBNSDBT. 108
IBNSDGA 108
IBNSDIG. 109
IBNSEXP. 110

© Copyright IBM Corp. 2004, 2007

IBNSFAN 110
IBNSFRU 111
IBNSLNK 111
IBNSNLS 113
IBNSPOW 113
IBNSREG 114
IBNSREM 115
IBNSSLC 116
IBNSSMM 118
IBNSSMR 118
IBNSSMU 119
IBNSSWE 119
IBNSSWP 120
IBNSTHM 120
IBNSVPD 121
Finding the other side of the link 122
Interpreting LEDs 122
Manipulating FRU identification LEDs 136
Isolating a problem with a port or link 136
Isolating performance problems 138
Suspected power problem 139
Suspected thermal problem 140
Verifying adapters are configured and available 140
Checking the subnet manager 140
Reordering Subnet Manager priority 142
Restarting the subnet manager 142
Synchronizing Subnet Manager time with HMC time 142
Location codes for machine types other than 7048 143
Checking switch software levels 144
Updating switch software 144
Switch is on the incorrect subnet 144
Understanding timestamp differences 145
Recovering from logical HCA configuration problems 145
Rebooting the entire switch chassis 145
Adjusting Firewall Parameters for SNMP Traps . . 146
GID-prefix procedures 146
Checking GID-prefixes 146
Setting GID prefixes 147
Logical identifier mask control procedures 147
Checking the logical identifier mask control . . 147
Setting the location identifier mask control . . 148
Repairs when using IBM Network Manager . . . 148
Plugging Rules for 7048-270 or SFS7008P switch line interface modules and fabric controllers . . . 152
Disappearing IBM Network Manager windows . . 152
Verifying Static12x or 4x configuration to a port 153
Determining faulty fabric controller cards versus faulty LIM cards 153
Missing LIM(s) or LIM ports 153

Administrative procedures for InfiniBand switches 157
Accessing a serial port on a switch 157
Accessing an InfiniBand switch command line interface 158
Reviewing the ts_log file 159
Setting IP addressing in switches 159
Setting IP addressing for a switch network with a single HMC 159
Setting IP addressing for a network with multiple HMCs 160
Rebooting switch cards 160
Setting database synchronization 161
Adjusting database-synchronization timeout . . . 161
Accessing FRU Identification LEDs 162
Filing bug reports 163

Diagnostics for networks managed by IBM Network Manager 165
Link diagnostic procedures 165
InfiniBand switch card diagnostics 173

Symbolic FRUs for InfiniBand cluster networks that are managed by the IBM Network Manager 175
How to find an adapter FRU using a reference code extension 176
How to find an adapter FRU using a reference code extension 177
CBLCONT 178
IBNACRD 178
IBNAPRT 178
IBNCCAB 179
IBNPBAT 181
IBNPFAN 182
IBNPPOW 182
IBNSCHS 183
IBNSCRD 184
IBNSCTL 185
IBNSPLN 186
IBNSSVC 186

Status procedures for the IBM Network Manager 189
Management properties view status procedure . . 189
End-Point Topology window status 190
Power status in End-Point Topology window 195
Logical topology window status 199
Switch topology window status 200
Switch Topology - Switch Properties: System Status 202
Switch Topology – Port Properties: System Status 203
Switch Environmental Status 205
Common Status Procedures 211

InfiniBand component location codes 217
Procedures for finding FRUs 218
Location codes used by the IBM Network Manager 218
How to find a switch FRU with a valid location code 222
How to find an adapter FRU with a valid location code 222
How to find a switch FRU using another device with a valid location code 223
How to find an adapter FRU using another device with a valid location code 223

How to find a switch FRU using a reference code extension 224
How to find an adapter FRU using a reference code extension 226
How to determine a switch FRU using the error description 227
How to determine an adapter FRU using the error description 227

InfiniBand parts information 229
How to determine an unknown part number for an InfiniBand switch network component 229
Identifying InfiniBand cables and cabling expansion units. 229
Removing and replacing InfiniBand parts 229
Parts catalog. 229

Appendix A. Notices 231
Trademarks 232
Regulatory notices. 233
Class A Notices - Federal Communications Commission (FCC) statement 233
Class B Notices - Federal Communications Commission (FCC) statement 235
Terms and conditions. 238
Product recycling and disposal 238
Battery return program 239
IBM Cryptographic Coprocessor Card Return Program 239

Appendix B. Planning and Installation Worksheets 241
Installation coordinating worksheets. 241
Planning checklist 242
Planning worksheets 242

Clustering systems using InfiniBand hardware

IBM® Systems server hardware supports clustering through Host Channel Adapters (HCAs) and InfiniBand switches.

This guide provides planning and installation information to help guide you through the process of installing a cluster fabric that incorporates InfiniBand switches. Information about how to manage and service an InfiniBand cluster is also included.

The following illustration shows servers that are connected in a cluster configuration with InfiniBand switch networks (fabric). The servers in these networks can be connected through switches with GX adapter HCAs.

Note:

1. In this information, *switch* refers to the InfiniBand technology switch unless otherwise noted.
2. Not all configurations support the following network configuration. Refer to your IBM sales information for supported configurations.

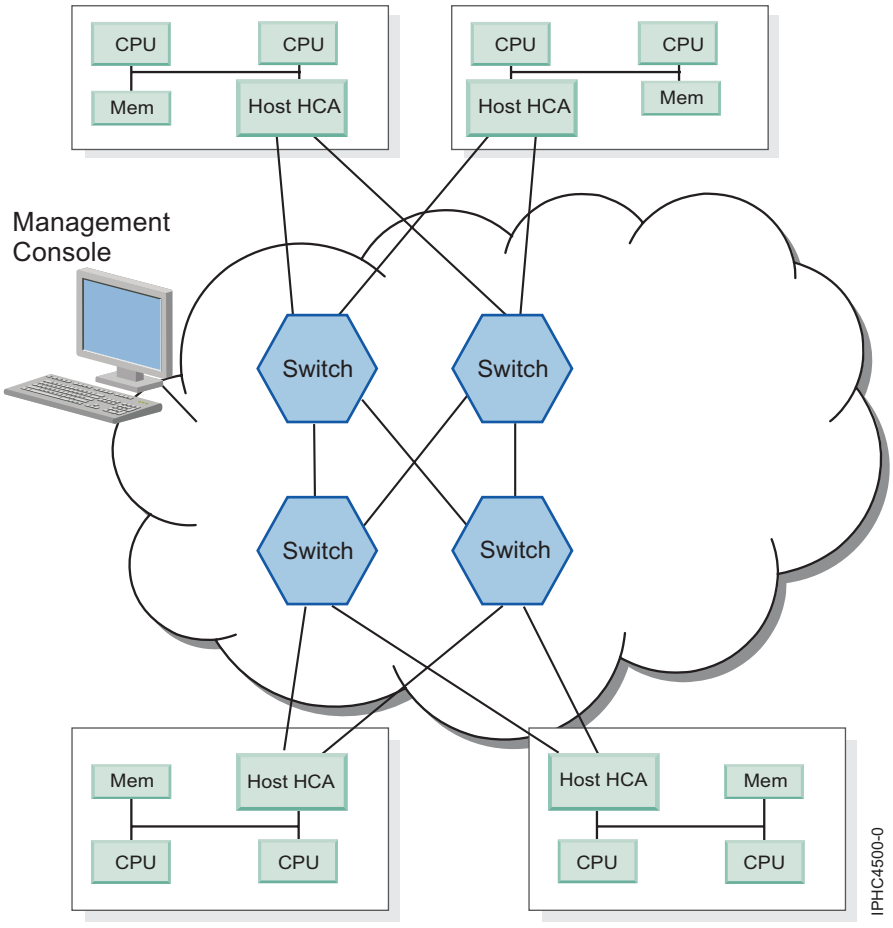

Figure 1. An InfiniBand network managed by a Hardware Management Console (HMC) with four switches (fabric) and four servers connected.

The following switch models are available through IBM:

© Copyright IBM Corp. 2004, 2007

- 7048-120 (Topspin 120 Server Switch)
- 7048-270 (Topspin 270 Server Switch)

The following switch models are available through Cisco:
- SFS7000P (Cisco SFS 7000 Server Switch)
- SFS7008P (Cisco SFS 7008 Server Switch)

If you are planning to use an InfiniBand switch in your clustered system configuration, the following table provides a list of InfiniBand hardware and software components available.

Table 1. InfiniBand network hardware solutions

Solution components	Specific InfiniBandhardware solution channel adapter
Adapter	GX 4x/12x Host Channel Adapter (HCA)
Systems	IBM System p5™ mid range and high-end
Switches	Topspin 120, Cisco 7000, Topspin 270, Cisco 7008
Cables	IBM certified cables
Fabric Management	IBM Network Manager running on an HMC
AIX version	AIX 5L™ Version 5.3 with the 5300-04 Technology Level plus APAR IY84727.
Linux version	SUSE Linux® Enterprise Server 9 (SP3) with IBM GX HCA driver and Open Fabrics Enterprise Distribution

Note: For the most recent information regarding cluster offerings, see the Facts and Features Web site: www.ibm.com/servers/eserver/clusters/hardware/factsfeatures.html.

Information resources

Documentation to help with planning, installing, managing, and servicing InfiniBand (IB) switch networks.

The following sections list sources of information that are grouped by planning, installing, and managing of an InfiniBand switch-network (fabric). The managing topics include administrative and service information.

The following documentation is available in the IBM Systems Hardware Information Center at http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp, the IBM eServer Cluster Information Center at http://publib.boulder.ibm.com/infocenter/clresctr/vxrx/index.jsp. Cisco and Topspin documentation is available from Cisco web site.

Planning for InfiniBand networks

The following lists documentation to help plan for an InfiniBand network:
- This information, *Clustering systems using InfiniBand hardware* , provides planning information to help guide you through installing a cluster fabric using InfiniBand switches. Go to http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/topic/iphau/infinbdpdf.pdf
- Managing your server using the Hardware Management Console (HMC) See http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/topic/ipha1/hardwaremanagementconsolehmc.htm.
- *Cluster Systems Management Planning and Installation Guide* for Cluster Ready Hardware Server. See the latest version at http://publib.boulder.ibm.com/infocenter/clresctr/vxrx/index.jsp?topic=/com.ibm.cluster.csm.doc/clusterbooks.html.
- *Topspin 120/Cisco SFS 7000 Hardware Guide*, order number: 10-00032-04-A0, See http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/topic/iphau/tpspn120hg.pdf.

- *Topspin 270/Cisco SFS 7008 Hardware Guide*, order number: 10-00044-04-A0, See http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/topic/iphau/tpspn270hg.pdf.
- *Cisco SFS 7000 Series Product Family Guides* for Cisco SFS7000P and SFS7008P. See http://www.cisco.com/en/US/products/ps6421/index.html. Under Technical Documentation & Tools, choose **Install and Upgrade**.
- Worldwide Customized Installation Instructions (WCII), service representative installation instructions, see http://w3.rchland.ibm.com/projects/WCII.
- IBM System p™ and AIX® Information Center Web site at

  http://publib.boulder.ibm.com/infocenter/pseries and http://publib.boulder.ibm.com/infocenter/pseries/v5r3/index.jsp?topic=/com.ibm.aix.doc/aixbman/commadmn/tcp_ip_over_infiniband.htm

Installing InfiniBand networks

The following guides are available to help you set up your network:

- This information, *Clustering systems using InfiniBand hardware* , provides information to help guide you through installing a cluster fabric using InfiniBand switches. Go to http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/topic/iphau/infinbdpdf.pdf
- *Managing your server using the Hardware Management Console (HMC)*. See http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/topic/ipha1/hardwaremanagementconsolehmc.htm.
- *Cluster Systems Management Planning and Installation Guide* (especially for Cluster Ready Hardware Server). See the latest version, http://publib.boulder.ibm.com/infocenter/clresctr/vxrx/index.jsp?topic=/com.ibm.cluster.csm.doc/clusterbooks.html.
- *InfiniBand Hardware Installation and Cabling Guide Web Release, Topspin number: 10-00122-WEB*
- *Topspin 120/Cisco SFS 7000 Quick Start Guide*, order number: 10-00033-04-A0
- *Topspin 270/Cisco SFS 7008 Quick Start Guide* , order number: 10-00045-04-A0
- *Cisco SFS 7000 Series Product Family Guides* for Cisco SFS7000P and SFS7008P. See http://www.cisco.com/en/US/products/ps6421/index.html. Under Technical Documentation & Tools, choose Install and Upgrade.

Administering and servicing InfiniBand networks

If you need to do administrative tasks on your InfiniBand cluster, the following guides are available to help you. This documentation is also intended for use by service personnel:

- This information, *Clustering systems using InfiniBand hardware* , provides information to help administer and service a cluster fabric using InfiniBand switches. Go to http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/topic/iphau/infinbdpdf.pdf
- *Managing your server using the Hardware Management Console (HMC)*. See http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/topic/ipha1/hardwaremanagementconsolehmc.htm.
- *IBM Cluster Systems Management for AIX 5L and Linux Planning and Installation Guide* (especially for Cluster Ready Hardware Server). See the latest version, http://publib.boulder.ibm.com/infocenter/clresctr/vxrx/index.jsp?topic=/com.ibm.cluster.csm.doc/clusterbooks.html.
- *Cisco SFS 7000 Series Product Family Guides* for Cisco SFS7000P and SFS7008P. See http://www.cisco.com/en/US/products/ps6421/index.html. Under Technical Documentation & Tools, choose *Maintain and Operate* or *Reference Guides*.

InfiniBand switch reference information

Each switch has a hardware guide that provides service information. Use the location code to help determine which guide you need.

To determine which guide to use when doing service tasks on a switch, use the location code for the switch unit and find the book in the following table.

Switch chassis	Unit location code	Manual
Topspin 120 Server Switch Cisco SFS 7000 Server Switch	U7048.120.*serial number* U7000P*sssss*.*serial number**	*Topspin 120/Cisco SFS 7000 Hardware Guide*, order number: 10-00032-04-A0
Topspin 270 Server Switch Cisco SFS 7008 Server Switch	U7048.270.*serial number* U7008P*sssss*.*serial number**	*Topspin 270/Cisco SFS 7008 Hardware Guide*, order number: 10-00044-04-A0
* *sssss* represents an arbitrary string of characters.		

Determining switch chassis from a location code

For switch components, you can determine the type of switch chassis by the unit location within the unit location code.

The format for the unit location code for the 7048-120and 7048-270switches is:
U*MachineType*.*Model*.001.*SerialNumber*.

The format for the unit location code for the Cisco 7000 series switches is:
U*ProductName*.*ProductSerialNumber*.

Topspin and Cisco switches are given model numbers as shown in the following list:

Table 2.

Name	Model number
Topspin 120	7048-120
Cisco 7000	SFS7000P
Topspin 270	7048-270
Cisco 7008	SFS7008P

Determining the server from a location code

For server components, you can determine the type of cage by the unit location and serial number within the unit location code: U*MachineType*.001.*SerialNumber*.

Planning for InfiniBand networks

Plan for a cluster that uses an InfiniBand network. Learn about key elements of the planning process, and organize your existing planning information.

When planning a cluster with an InfiniBand network, you bring together many different devices and management tools to form a cluster. Major components to consider are:
- Servers
- I/O devices
- InfiniBand network devices
- Frames (racks)
- Service network, including:
  - Hardware managment consoles (HMC)
  - Ethernet devices
  - Cluster Systems Management (CSM) server (for multiple HMC environments)
  - AIX Network Installation Management (NIM) server (for servers with no removable media)
  - Linux distribution server (for servers with no removable media)
- System management applications, including:
  - HMC
  - CSM
  - IBM Network Manager
- Physical characteristics such as weight and dimensions
- Electrical characteristics
- Cooling characteristics

"Planning for InfiniBand networks" on page 2 is intended to help you find references for many required documents and other Internet resources that will help you plan your cluster. It is not an exhaustive list of the documents that you might need, but it provides a good starting point for gathering required information.

Use the Planning Overview section as a roadmap through the planning procedures. Read through the Planning Overview once without traversing any of the links to other sections so that you can understand the overall planning strategy. Then, go back and follow the links that direct you through the different procedures.

The Planning Overview, the planning procedures and the sub-procedures are arranged under headings in a sequential order for a new cluster installation. If you are performing an installation other than that for a new cluster, you may need to pick and choose which procedures to use, but you should do so in the order in which they appear in the Planning Overview. For a new cluster install, you can read straight through to the end of the *Planning for InfiniBand networks* section. If you are using the links in the *Planning Overview* section, you will want to note when a planning procedure ends so that you know when to return to the *Planning Overview*. The end of each major planning procedure will be indicated by [*planning procedure name*] ends here.

Planning overview

Do the following to start planning your cluster:

1. When planning your cluster, you need to first gather and review the planning and installation information for the various components in the cluster. Because this document provides supplemental information with respect to clustered computing with an InfiniBand network, you must understand all of the planning information for the individual components before proceeding with this planning overview.

© Copyright IBM Corp. 2004, 2007

2. Review the "Planning checklist" on page 22, which can help you keep track of which planning steps that you have completed.
3. Review the planning resources for the individual servers that you want to use in your cluster. This document will only address server planning with respect to Host Channel Adapters (HCAs). Some information from this planning exercise should be recorded in the "Server planning worksheet" on page 27. In addition to planning for your servers, you will need to plan for any frames that will contain the servers. Some information from this planning exercise should be recorded in the "Frame and rack planning worksheet" on page 26.
4. Review the information for HMC planning and Cluster Systems Management (CSM) planning so that you understand the layout of the service network and know the Ethernet devices that are required to support the service network. The following points must be considered when planning the service network:
   - With more than one HMC, the CSM planning information is more critical because you must plan for the Cluster Ready Hardware Server (CRHS) component in CSM. CSM must be enabled to manage discovery of all of the devices in the cluster, and to assign them to the controlling HMCs. If you have only a single HMC, you might choose to use CRHS with CSM.
   - When you have more than one HMC, or have chosen to use CRHS with CSM, the CSM server is required to be the DHCP server for the service network. For optimum performance, set up the CSM server as a standalone server. If you use one of the computing servers or I/O servers in the cluster for CSM, the CSM operation might degrade performance for user applications, and it will complicate the installation process with respect to server setup and discovery on the service network.
   - If you do not require a CSM server, you might need a server to act as an AIX Network Installation Management (NIM) server for eServer standalone diagnostics. This situation applies to servers that do not have removable media (CD or DVD), such as an IBM eServer™ p5575 (9118-575).
   - If you have servers with no removable media and they are running Linux partitions, you may require a server to act as a Linux distribution server.
   - If you require both an AIX NIM server and a Linux distribution server, and you choose the same server for both, a reboot is needed to change between the services. If the AIX NIM server is used only for eServer diagnostics, this might be acceptable in your environment. However, you should understand that this might prolong a service call if use of the AIX NIM service is required. For example, you might usually have the server acting as a Linux distribution server. If AIX NIM services are required for eServer standalone diagnostics during a service call, the server must be rebooted to AIX and then diagnostics can be performed.
5. To understand the minimal level of software and firmware required to support clustering with an InfiniBand network, review "Required levels of support, firmware, and devices" on page 7.
6. Having determined the types and numbers of components that you require in your cluster, plan your InfiniBand network cabling.
7. After planning your InfiniBand network cabling, you should plan your switch and HCA configuration settings. For switches see "Planning InfiniBand switch configuration" on page 8. For HCAs, see "Planning an IBM GX HCA configuration" on page 9.
8. Review the information about "Management subsystem planning" on page 10.
9. After you understand the basic concepts for cabling an InfiniBand network, review the high-level install flow information in "Installation flow of clusters using the IBM Network Manager" on page 11. There are hints about planning your InfiniBand network, and also guidelines to help you to coordinate between your (customer) and IBM Service Representative responsibilities during the installation process.
10. Consider special circumstances such as whether you are configuring a cluster for High Performance Computing (HPC) Message Passing Interface (MPI) applications, in which case, you should refer to "Planning for a high-performance computing message-passing interface configuration" on page 18.

11. If you are using 12x HCAs (for example in a 9119-590 server), you should review “Planning octopus cables in static 12x cabling” on page 19, to understand the unique cabling and configuration requirements when using these adapters with the available 4x switches.
12. For some more hints and tips on planning your network, review “Planning Aids” on page 22.

If you have completed all the previous steps, you can plan your network in more detail using the planning worksheets provided in “Planning worksheets” on page 23.

When you are ready to install the hardware from which you will build your cluster, review any information in READMEs and online information related to the software and firmware to ensure that you have the latest information and the latest supported levels of firmware.

Required levels of support, firmware, and devices

The following tables provide the minimum requirements necessary to support InfiniBand network clustering.

Note: For the most recent updates to this information, see the following Facts and Features Web site, http://www.ibm.com/servers/eserver/clusters/hardware/factsfeatures.html

Table 3. Verified and approved hardware associated with an IBM System p5 and eServer p5 InfiniBand cluster

Supported cluster component	Component model or feature code number
Servers	9119-590 9119-595 9118-575* 9117-570 9116-561 9131-52A 9113-550 9133-55A*
Switches	7048-120 7048-270 Cisco 7000 Server Switch** Cisco 7008 Server Switch**
Host Channel Adapters (HCAs)	The feature code is dependent on server. Order one or more GX, Dual-port HCA per server that requires connectivity to InfiniBand networks. The maximum number of HCAs allowed is dependent on the server model. See “Planning for InfiniBand networks” on page 5.
Note: * Validated to work in an IBM High Performance Computing (HPC) cluster. ** For approved IBM System p5 and eServer p5 InfiniBand configurations, see the Facts and Features Web site, http://www.ibm.com/servers/eserver/clusters/hardware/factsfeatures.html	

Table 4. Minimum levels of software and firmware associated with an InfiniBand cluster

Software	Minimum Level
AIX	AIX 5L Version 5.3 with the 5300-04 Technology Level plus APAR IY84727.

Table 4. Minimum levels of software and firmware associated with an InfiniBand cluster (continued)

Software	Minimum Level
SUSE Linux Enterprise Server 9	SUSE Linux Enterprise Server 9 (SP3) with IBM GX HCA driver and Open Fabrics Enterprise Distribution
Hardware Management Console	SQ7_0622A_0517
System firmware level for System p5	01SF235_160_160
InfiniBand switch software	Cisco release 2.5.0/build build 251 for a: • 7048-120 • 7048-270 • SFS7000P • SFS7008P

Required levels of support, firmware, and devices that are needed to support InfiniBand clusters ends here.

Planning InfiniBand network cabling and configuration

Use all the available resources when planning InfiniBandnetwork cabling and configurations.

Before you plan your InfiniBand network cabling, review the hardware installation and cabling information for your Cisco switch at http://www.cisco.com/en/US/products/ps6418/index.html. Search for the documentation regarding the SFS7000P or SFS7008P Series InfiniBand Server Switches. Review any documentation regarding installation and cabling of your switch.

While planning your cabling, keep in mind the IBM server and frame physical characteristics that affect cable planning. In particular, consider the following:

- server height. IBM System p5 and eServer p5 servers, when used in an IBM System p5 cluster, are either 2U or 4U in height.
- routing to the cable entrance of a frame
- routing within a frame
- floor depth

If you are using 12x HCAs (for example in a 9119-590 server), review “Planning octopus cables in static 12x cabling” on page 19, to understand the unique cabling and configuration requirements when using these adapters with the available 4x switches.

Use the “Switch planning worksheet” on page 28 to record the switch port connections and use the “Server planning worksheet” on page 27 to record the HCA port connections.

Planning InfiniBand network cabling and configuration ends here.

Planning InfiniBand switch configuration

InfiniBand switches require some custom configuration to work correctly in an IBM System p5 cluster. The configuration settings that need to be planned are:

- IP addressing on the service Ethernet network configured as DHCP or static. If addressing is static, then assign the address.
- GID-prefix to identify the subnet of which the switch is a member.
- LMC value
- Switch name
- Any 12x cabling considerations

You must determine if the IP addressing a switch will have on the service Ethernet network is configured for DHCP or static addressing. A 7048-270 or SFS7008P might have two I/O management modules and thus require a single address with two connections on the same service Ethernet network. For details, see "Management subsystem planning" on page 10.

Each subnet in the InfiniBand network must be assigned a GID prefix, which will be used to identify the subnet for addressing purposes, and within IBM Network Manager. The GID-prefix is an arbitrary assignment with a format of: *xx:xx:xx:xx:xx:xx:xx:xx*. For example: `FE:80:00:00:00:00:00:01`.

If you are using a switch in an IBM HPC cluster, you must set the LMC value to 2. For more details, see "Planning for a high-performance computing message-passing interface configuration" on page 18. Assign each switch a name. This is an arbitrary assignment, but it is useful to have a naming convention that includes an indicator of which frame the switch is in, and which device it is within the frame. For example, `ezfr11sw2` indicates that this is in frame 11 (fr11) and it is the second switch from the bottom (sw2), where `ez` is being used as a base name for the entire cluster (`ez = example cluster`).

Finally, if this is a 4x switch connecting to 12x HCA, you will require octopus cables to connect (3) 4x switch ports to the 12x HCA. For more details, see "Planning octopus cables in static 12x cabling" on page 19.

Use the "Switch planning worksheet" on page 28 to record switch planning information.

Planning InfiniBand switch configuration ends here.

Planning an IBM GX HCA configuration

An IBM GX Host Channel Adapter (HCA) needs to have certain configuration settings to work in an IBM System p5 InfiniBand cluster:

- GUID index
- Capability
- GID-prefix for each port of an HCA

Each physical HCA contains a set of 64 Globally Unique IDs (GUIDs) that can be assigned to partition profiles. These GUIDs are used to address Logical HCA (LHCA) resources on an HCA. You can assign multiple GUIDs to each profile, but you can assign only one GUID from each HCA to each partition profile. Each GUID can be used by only one logical partition at a time. You can create multiple partition profiles with the same GUID, but only one of those partition profiles can be activated at a time.

The GUID index is used to choose one of the 64 GUIDs available for an HCA.. The GUID can be any number from 1 through 64. You will usually assign a GUID index based on which LPAR and profile you are configuring. For example, on each server you might have 4 partitions, and the first one on each server might use a GUID index of 1, and the second one on each server would use a GUID index of 2, right through to the fourth using a GUID index of 4.

The Capability setting is used to indicate the level of sharing to be done and can be one of the following:

Low A low setting results in allocation of 1/16 of the adapter resources

Medium
A medium setting results in allocation of 1/8 of the adapter resources

High A high setting results in allocation of 1/4 of the adapter resources

Dedicated
A dedicated setting results in allocation of all of the adapter resources

While the GID-prefix for a port is not something that you explicitly set, it is important to understand the subnet to which a port attaches. This is determined by the switch to which the HCA port is connected. The GID-prefix is actually configured for the switch. See "Planning InfiniBand switch configuration" on page 8.

Note: When a Host Channel Adapter (HCA) is added to a logical partition, the HCA becomes a required resource for the partition. If the HCA ever fails in such a way that the system's GARD function prevents it from being used, the logical partition cannot be reactivated. If this occurs, a pop-up message displays on the controlling HMC which indicates that you need to unassign the HCA from the logical partition to continue activation. The GARD function is invoked for serious adapter or bus failures that could impair system operation, such as ECC errors or state machine errors. InfiniBand link errors should not invoke the GARD function.

Additional information on partition profiles is available in the Information Center: http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/topic/iphat/iphatparprofile.htm

Use the "Server planning worksheet" on page 27 to record the HCA configuration settings.

Planning an IBM GX HCA configuration ends here.

Management subsystem planning

The management subsystem comprizes the Ethernet service network, Hardware Management Consoles (HMCs), Cluster Systems Management (CSM) Management Server, AIX Network Installation Management (NIM) server, and Linux Distribution server. There are links to key references in planning the management subsystem.

A customer-supplied Ethernet service network is required to support the InfiniBand cluster computing environment. The number of Ethernet connections depends on the number of servers, Bulk Power Controllers (BPCs) in 24-inch frames, InfiniBand switches, and HMCs in the cluster. The CSM Management Server with use of the Cluster-Ready Hardware Server (if applicable) also requires a connection to the Ethernet service network.

Note: While you might have two service networks on different subnetworks to support redundancy in IBM servers, BPCs, and HMCs. The InfiniBand switches support only a single service network (even though some InfiniBand switch models have multiple Ethernet connections).

An HMC is required to manage the LPARs and to configure the GX bus HCAs in the servers, as well as to run IBM Network Manager. The maximum number of servers that can be managed by an HMC is 32. When you exceed 32 servers, additional HMCs are required. See Solutions with the Hardware Management Console in the IBM Systems Hardware Information Center, http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp. Click **Planning > Solution planning > Planning for consoles, interfaces, and terminals**.

If you have a single HMC in the cluster, it is usually configured to be the required DHCP server for the Ethernet service network. If a CSM and a Cluster Ready Hardware Server are being used in the cluster, the CSM Management Server must be configured as the DHCP server for the Ethernet service network.

If you require more than one HMC to manage your cluster's servers and switches, you must use CSM on a CSM Management Server, and you must configure a Cluster-Ready Hardware Server as the DHCP sever on the Ethernet service network. Refer to the *IBM Cluster Systems Management for AIX 5L and Linux Planning and Installation Guide*.

Because there can be only one DHCP server per Ethernet service network, if you add a CSM Management Server it must be configured as the DHCP sever on the Ethernet service network, and you must reset the HMC that was previously the DHCP server to a fixed IP address for the network. Then the

Cluster-Ready Hardware Server must be configured to recognize the cluster's servers, BPCs, HMCs and InfiniBand switches. Refer to the *IBM Cluster Systems Management for AIX 5L and Linux Planning and Installation Guide*.

The servers have connections to the Ethernet service network. For more information about connecting to the Ethernet service network, refer to the server documentation. In particular, pay attention to:

- The number of service processor connections from the server to the service network
- If there is a Bulk Power Controller (BPC) for the power distribution, as in a 24-inch frame, the BPC can provide a hub for the system units in the frame, and thus allow for a single connection for each frame to the Ethernet service network.

After you understand the number of devices and cabling of your service Ethernet network, you will need to consider the device IP addressing. Do the following:

1. Determine the domain addressing and netmasks for the one or two Ethernet service networks that you will implement.
2. Assign static-IP addresses:
   - You need to assign a static IP address for HMCs when you are using CSM and Cluster-Ready Hardware Server. This is required when you have multiple HMCs in the cluster.
   - You need to assign a static IP address for switches when you are using CSM and Cluster-Ready Hardware Server. This is required when you have multiple HMCs in the cluster.
3. Determine the DHCP range for each subnet.
4. If you must use, or choose to use CSM and Cluster-Ready Hardware Server, the DHCP server is recommended to be on the CSM Management Server, and all HMCs must have their DHCP server capability disabled. Otherwise, you are in a single HMC environment where the HMC is the DHCP server for the Ethernet service network.

If there are servers in the cluster without removable media (CD or DVD), an AIX NIM server for eServer diagnostics is required. If you are using the AIX operating system in any of your partitions, this provides NIM service for the partitions.

If there are partitions running Linux in your servers, and they do not have removable media (CD or DVD), you will require a Linux Distribution server. Use the "Cluster summary worksheet" on page 25 to record the information in this section.

Management subsystem planning ends here.

Installation flow of clusters using the IBM Network Manager

The high-level flow (or sequence of events) needed to install a cluster with IBM Network Manager support.

Review the following items:

1. "Key Installation points"
2. "Installation responsibilities by organization" on page 12
3. "Installation responsibilities by units and devices" on page 13
4. "Order of Installation" on page 13
5. "Installation Coordination Worksheets" on page 18

Key Installation points

When you are coordinating the installation of the many systems, networks and devices in a cluster, several factors can influence the success of the installation:

1. The order of arrival of physical units is important. While units may be placed physically on the floor in any order after the site is ready for them, there is a definite order in how they are to be powered and recognized on the Ethernet service network, as well as how you want to cable them for the best possible cable management.
2. The types of units and contractual agreements can affect the composition of the installation team. The team can consist of customer, IBM, or vendor personnel. For more guidance about installation responsibilities, see "Installation responsibilities by organization."
3. The level of planning done for the cluster can significantly affect how well the actual installation is coordinated and completed.
4. If you have 12x HCAs and 4x switches, switches must be powered on and configured with proper 12x groupings before servers are powered on. The order of port configuration on 4x switches configured with groups of 3 ports acting as a 12x link is not consistent. Therefore, specific steps must be followed to ensure that the 12x HCA is connected as a 12x link and not a 4x link.
5. Because of the addressing methods used by InfiniBand switches, all switches in a cluster network must be connected to the same Ethernet service network. If there are redundant connections available on a switch, they must also be connected to the same Ethernet service network.

Installation responsibilities by organization

Within a cluster that has an InfiniBand network, different organizations are responsible for various installation activities. However, it is possible for the specific responsibilities to change because of agreements between the customer and the supporting hardware teams.

Note: Given the complexity of typical cluster installations, a trained, authorized installer should perform the installation.

Customer installation responsibilities:

- Set up of management consoles (HMCs and CSM management servers)
- Install customer setup units (servers)
- Update system firmware
- Any updates or customization that are performed using IBM Network Manager:
  - Update InfiniBand switch software using IBM Network Manager
  - Customize InfiniBand network configuration using IBM Network Manager
  - Customize HCA partitioning and configuration
- Any verification that is performed using IBM Network Manager or IBM Service Focal Point:
  - Verifying the discovery of switches and HCAs by IBM Network Manager
  - Verifying the InfiniBand network topology and operation

IBM installation responsibilities:

- Install and service IBM systems (servers that are not customer installable) and HCAs. These include the System p5 570, 575, 590 and 595 systems.
- Verify server operation for IBM installable units
- Service IBM supplied InfiniBand cables, but not their installation.

Vendor installation responsibilities:

Note: This document cannot detail the contractual possibilities for vendor responsibilities. By contract, the customer might be responsible for some of these activities.

- Install switches
- Set up the Ethernet service network IP and attach switches to the service network
- Cable the InfiniBand network

Installation responsibilities by units and devices

Note: It is possible that contracted agreements might alter the basic installation responsibilities for particular devices.

Installation responsibilities for servers:
The use of a server in a cluster with an InfiniBand network does not change the usual installation and service responsibilities for it. Some servers are installed by IBM and others are installed by the customer. See the specific server documentation to determine who is responsible for the installation.

Installation responsibilities for HMCs:
The type of servers attached to the HMCs dictate who installs them. To determine who is responsible for the installing the HMC, see the HMC documentation. This is usually the responsibility of the customer or IBM service personnel.

Installation responsibilities for CSM:
CSM is needed in a multiple HMC environment so that there is a centralized source for device discovery in the cluster. The customer is responsible for CSM installation.

Installation responsibilities for InfiniBand switches:
The switch manufacturer or its designee (business partner) or another contracted organization determines who is responsible for installing the switches.

Installation responsibilities for switch network cabling:
The customer must work with the switch manufacturer or its designee or another contracted organization to determine who is responsible for installing the switch network cabling. However, if there is a failure in a network cable with an IBM part number, IBM service is responsible for servicing the cable.

Installation responsibilities for Ethernet service network devices:
Any Ethernet devices such as switches or routers that are required for the service network are the responsibility of the customer.

Installation responsibilities for service network cabling:
The organization that is responsible for installing a device is responsible for connecting the device to the service network.

Order of Installation

This section provides a high-level outline of the general tasks required for installing a new cluster. If you understand the full installation flow for installing a new cluster, you can more easily identify the tasks that you will perform when you expand your InfiniBand cluster network. Tasks, such as adding InfiniBand hardware to an existing cluster, adding Host Channel Adapters (HCAs) to an existing InfiniBand network, and adding a subnet to an existing network. These additional installation tasks are discussed later in this section. To complete a cluster installation, all devices and units must be available before you begin to install the cluster. Fundamental tasks for installing a cluster include:

1. Ensuring that the site is set up with power, cooling, and floor requirements.
2. Installing and configuring switches and processing units.
3. Connecting cabling of units to the service network.
4. Verifying that the units can be discovered on the service network.
5. Verifying basic unit operation.
6. Connecting cabling for the InfiniBand network.
7. Verifying the InfiniBand network topology and operation

The following figure separates the tasks by major subsystem. The following list illustrates the preferred order of installation by major subsystem. The order minimizes potential problems with having to perform recovery actions while you install the cluster, and also minimizes the number of reboots of devices during the installation.

1. Management consoles and the service network. Management consoles include the HMC, and any server running CSM.
2. Servers in the cluster
3. Switches
4. Switch cable installation

By breaking down the installing tasks by major subsystem, you can see how to install the units in parallel, or how you might be able to perform some installation tasks for on-site units while waiting for other units to be delivered. It is important that you recognize the key points in the installation where you cannot proceed with one subsystem's installation task before completing the installation tasks in the other subsystem. These are called *merge* points, and are illustrated using the following symbol:

Figure 2. Merge points symbols

Some key merge points are:

1. The Management Consoles must be installed and configured before starting to cable the service network. This will allow proper DHCP management of the IP addressing on the service network. Otherwise, the addressing may be compromised.
2. You must power on the InfiniBand switches and configure their IP addressing before connecting them to the service network.
3. Before attaching the cables to the HCAs in servers that have been connected to a power source, if you have 12x HCAs to connect to 4x switches, you must power on the switches, connect the cables to their ports, and configure the 12x groupings. This is so that the auto-negotiation to 12x by the HMCs can occur correctly. When powering up the switches, it is not guaranteed that the ports will activate in an order that will make the link appear as 12x to the HCA. Therefore, you must be sure that the switch is properly cabled, configured, and ready to negotiate to 12x before starting the servers that contain HCAs.
4. To fully verify the InfiniBand network, the servers must be fully installed in order to pass data and run any tools required to verify the network. You can verify the topology before this by using IBM Network Manager and checking the topology for neighboring servers. However, the servers must be powered on to Standby to perform this verification.

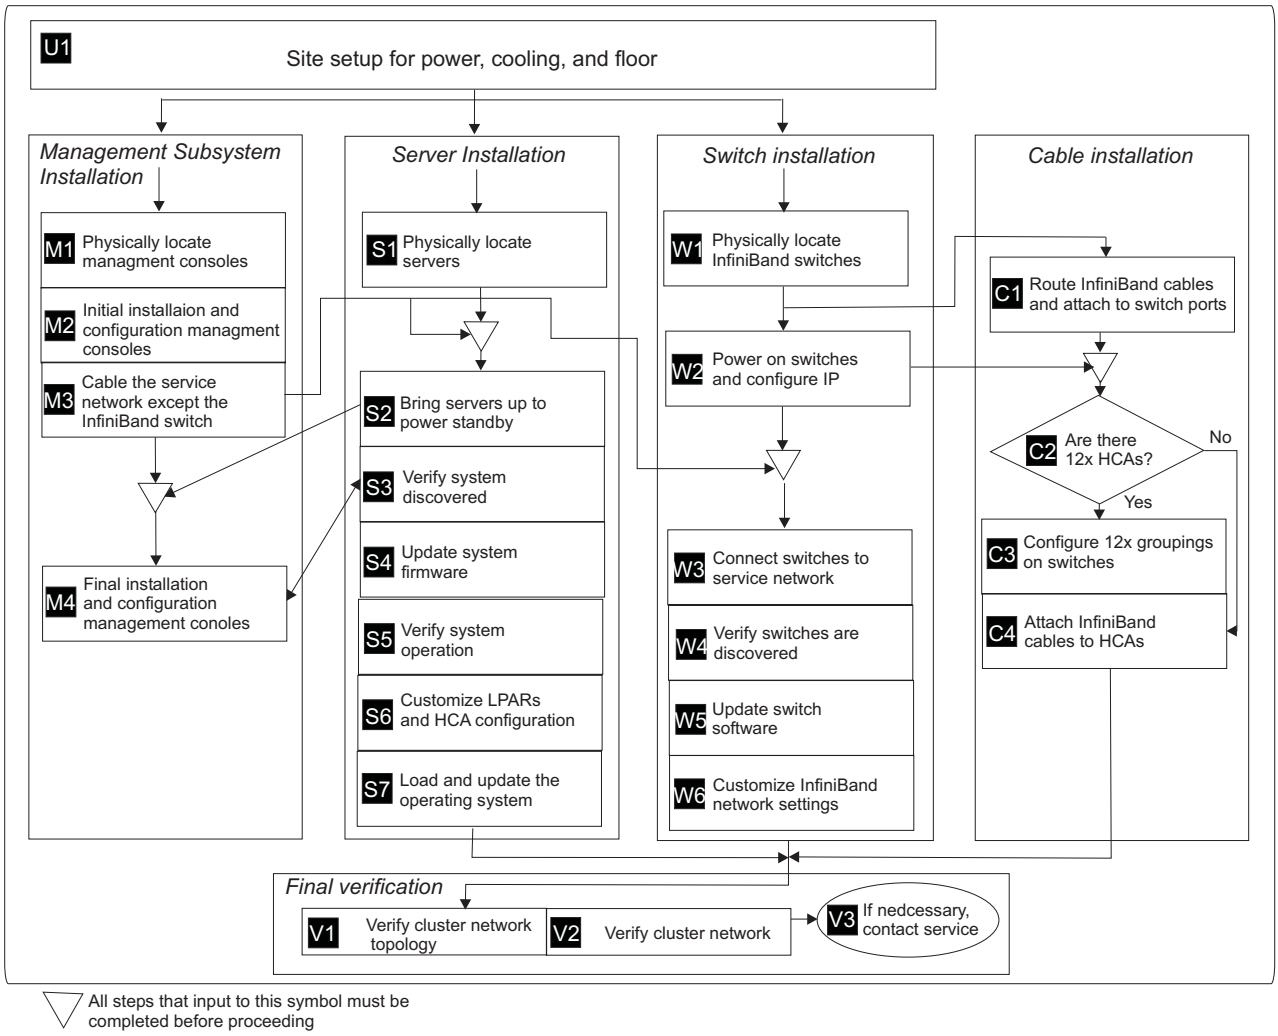

Figure 3. High-level cluster installation flow

Important: In each task box of the preceding figure, there is also an index letter and number. These indexes indicate the major subsystem installation tasks and you can use them to cross-reference between the following descriptions and the tasks in the figure.

The tasks indexes are listed before each of the following major subsystem installation items:

U1 Set up the site for power, cooling and if applicable, prepare floor cutouts for cable routing.

M1, **S1**, **W1**
Place units and frames in their correct positions on the floor. This includes, but is not limited to HMCs, CSM management servers and cluster servers (with HCAs, I/O devices, and storage devices) and InfiniBand switches. You can physically place units on the floor as they arrive, however do not apply power nor cable any units to the service network or to the InfiniBand network until instructed to do so.

M2 Do the initial management console installation and configuration. This includes HMCs CSM console, and DHCP service for the service network.

Note: If these devices and associated services are not set up correctly before applying power to the base servers and devices, you will not be able to correctly configure and control cluster devices. Furthermore, if this is done out of sequence, the recovery procedures for doing

this part of the cluster installation can be quite lengthy.
Set up static addresses for HMCs and InfiniBand switches. Set up the DHCP ranges.

M3 Connect service processors and Bulk Power Adapters (BPAs) to the service network. Do not attach InfiniBand switches to the Ethernet service network at this time.

Note: Switch IP addressing does not default as a DHCP client. In certain configurations, the switch's service network IP addressing will be static. In other configurations, the switch's service network IP addressing must be configured as a DHCP client.

M4 Do the final management console installation and configuration which involves assigning or acquiring cluster servers to their managing HMCs and authenticating frames and servers through the cluster-ready hardware server. This is not required when you are not using CSM and a cluster-ready hardware server.

Note:

1. The double arrow between M4 and S3 indicates that these two tasks cannot be completed independently. As the server installation portion of the flow is completed, then the management console configuration can be completed.
2. When a cluster requires multiple HMCs, CSM is required to help manage device discovery. In this case, the setup of CSM and the cluster-ready hardware server peer domains is critical to achieving correct cluster device discovery. It is also important to have a central DHCP server, which is recommended to be on the same server as CSM.

When this process is complete, the BPAs and cluster server's service processors must be at power standby state. To be at the power standby state, the power cables for each server must be connected to the appropriate power source, Prerequisites for M4 are M3 and S2 . Corequisite for M4 is S3 .

The following server installation and configuration operations (S2 through S7) and M3 can be done simultaneously. Do the following:

M3 Attach the cluster server's service processors and BPAs to the service network. This must be done before connecting power to the servers, and after the management consoles are configured so that the cluster servers can be discovered correctly.

S2 To bring the cluster servers to the power standby state, connect the cluster's servers to their appropriate power sources. Prerequisites for S2 are M3 and S1

S3 Verify the discovery of the cluster servers by the management consoles.

S4 Update the system firmware.

S5 Verify the system operation.

S6 Customize LPAR and HCA configurations.

S7 Load and update the operating system.

Do the following Switch installation and configuration tasks W2 and W3 :

W2 Power on and IP configuration of the switch Ethernet connections. This must be done before attaching it to the service network.

W3 Connect switch to the service networks. If there is more than one network, all switches must be attached to a single service network, and all redundant switch Ethernet connections must be attached to the same network. Prerequisites for W3 : M3 W2

W4 Verify discovery of the switches by IBM Network Manager.

W5 Update switch software using IBM Network Manager.

W6 Customize InfiniBand network configuration

Do the following to cabling the InfiniBand network

Note:

1. It is possible to cable and start networks other than the InfiniBand networks before cabling and starting the InfiniBand network.
2. When plugging InfiniBand cables between switches and HCAs, connect the cable to the switch end first. This is particularly important in this phase of the installation.

C1 Route cables and attach cables ends to the switch ports. Apply labels at this time.

C2 **C3**
If 12x HCAs are connecting to 4x switches, the switch ports must be configured and the switches must be powered on before attaching the cables to the HCA ports. Prerequisites for **C2** are **W2** and **C1**

C4 Attach the InfiniBand cable ends to the HCA ports.

Do the following to verify the cluster networking topology and operation:

V1 This will involve checking the topology by using IBM Network Manager. A vendor might have an alternative method for checking the topology. Prerequisites for **V1** are **S7** , **W6** and **C4** . It is possible to start this after **S5** the servers have booted to power standby. However, you cannot proceed to **V2** until after **S7** is completed.

V2 You must also check for serviceable events reported to Service Focal Point. Furthermore, an all-to-all ping is suggested to exercise the InfiniBand network before putting the cluster into operation. A vendor may have an alternative method for verifying network operation. However, Service Focal Point should always be consulted, and serviceable events should be addressed. If a vendor has discovered and resolved a serviceable event, then the serviceable event should be closed out of Service Focal Point.

Pre-requisites for **V2** : **V1** **S7** **W6** **C4**

V3 You might have to contact service numbers to resolve problems after service representatives leave the site.

Installation Coordination Worksheets

Use the following worksheets to help coordinate installation tasks. Each organization should use a separate installation sheet and it should be completed using the flow illustrated in Figure 3 on page 15.

Table 5. Sample Installation coordination worksheet

Organization:				
Task	**Task Description**	**Prerequisite tasks**	**Scheduled date**	**Completed date**

Table 6. Example installation coordination worksheet

Organization: IBM Service				
Task	**Task Description**	**Prerequisite tasks**	**Scheduled date**	**Completed date**
S1	Place the model 575s on floor		8/18/2006	
M3	Cable the model 575s and BPAs to service network			
S2	Bring up the model 575s			
S3	Verify discovery of the system and system operation			
S5	Verify System operation			

Installation flow of clusters using the IBM Network Manager ends here.

Planning for a high-performance computing message-passing interface configuration

Assumptions:

- Proven configurations for High Performance Computing (HPC) Message Passing Interface (MPI) are limited to:
  - One adapter per Operating System Image (OSI) with one link connected to the switch (that is one subnet or plane).
  - One adapter per OSI with two links connected to independent switches (that is more than one subnet or plane).
  - Two adapters per OSI with one link off each adapter connected to independent switches (that is more than one subnet or plane).

Limitations:

While most aspects of a cluster are not affected by the use of IBM Network Manager, redundant HMCs are not supported.

Other considerations:

Because HPC applications are designed particularly with performance in mind, it is important to configure the InfiniBand network components with this in mind. The main consideration is that the LID Mask Control (LMC) field in the switches must be set to provide more LIDs per port than the default of one. This provides more addressability and better opportunity for using available bandwidth in the network. The HPC software provided by IBM is designed to provide optimum performance with an LMC value of 2. The number of LIDs is equal to 2^x, where x is the LMC value. Therefore, the LMC value of 2 that is required for IBM HPC applications will result in four (4) LIDs per port.

The LMC setting planned here should be recorded in a "Switch planning worksheet" on page 28 which is meant to record switch configuration information.

Important information for planning an high performance computing MPI configuration ends here.

Planning octopus cables in static 12x cabling

Because InfiniBand GX Host Channel Adapters (HCAs) operate at a 12x speed, you must use octopus cables to connect to a group of three 4x ports on the InfiniBand switch.

Overview of octopus cables in static-12x configurations

Octopus cables have a 12x connector on one end and three 4x connectors on the other end. This allows a 12x device to connect to a group of three switch ports to enable full 12x bandwidth (30Gbps) where the three 4x switch ports and the 12x device port operate as a single 12x link, or interface. The grouping of three 4x ports is often referred to as a cluster of ports, or simply a cluster. To avoid confusion with the concept of clustering servers, when referring to groups of 4x switch ports, the term *group* is used instead of cluster.

When a group of three ports is configured in static-12x mode, the ports operate in unison with a single port assuming control of the group. This controlling port will be the port that you will see as Active, and reporting an operating speed of 30Gbps. Furthermore, static-12x groups only operate at a speed of 12x. They will not degrade to 4x or 1x speeds. Either the link is running at 12x with the controlling port being Active or the link is Down and no data can be passed.

The 4x port ends of the cable are labeled 16, 17 and 18. This information does not address 12x to 12x configurations, nor does it address auto-negotiate-12x mode.

Behaviors of octopus cable connections in static-12x configurations

The most significant behaviors of static-12x groups are:

- The lowest numbered port in a group is the *configuration* port. It determines if the group is running in static-12x mode. The lowest numbered port in a group is the only port that can be enabled or disabled while in static-12x mode. The other two ports' configuration properties (admin-speed, admin-negotiate and admin-status) are irrelevant.
- The controlling port in a group can be different for each group, and it may be different from the lowest numbered, or configuration port. It is either the lowest numbered port or the highest numbered port. It will always be connected to the 16 connector of the Octopus cable. Cabling diagrams are found in "Planning 4x connections for octopus cables" on page 21.
- The controlling port is the only active port of the three ports in the group. When looking at the port status, or *oper-status* of the three ports, the controlling port is *active* and the other two ports will be *down*.
- With a 7048-120 or SFS7000P, the LED for the controlling port that is active might be solid green or blinking green. The other two ports will be solid green.
- With a 7048-270 or SFS7008P, when a group's status is active the traffic LED for the controlling port in the group is flashing and the traffic LEDs for the other two, non-controlling ports are solid green. The logical LED for the controlling port is on, but logical LEDs for the other two, non-controlling, ports are off. For each of the non-controlling ports in the group, when the traffic LED is on and the logical LED is off, their physical link is established but their logical link is not established. The logical link for the group is established on the controlling port of the group and this is indicated by the logical LED of the controlling port.
- If you shift the group of three 4x connectors to the left or to the right one or more ports, it is possible to establish a connection on one or more of the ports adjacent to the group, if the adjacent group of

ports is not configured to static-12x mode. However, you will not see all three green LEDs come on unless all three connectors are outside of the intended group.

- If a connection to any of the three 4x ports fails, the entire link will go down, and all three green LEDs will go off.

Planning octopus cable connections in static-12x configurations

1. Determine which groups will be used for static-12x mode. See "Determining groups to use for static-12x mode, 7048-120 or SFS7000P."
2. Configure the groups for static-12x mode, by configuring the lowest numbered port in the group to static-12x mode. See "Configuring Static-12x groups."
3. Connect the three 4x connectors of the cable to the appropriate switch ports. See "Planning 4x connections for octopus cables" on page 21.

Determining groups to use for static-12x mode, 7048-120 or SFS7000P

For a 7048-120 or SFS7000P, each 4x switch has 24 ports in two rows. These ports are logically divided into groups of three, as shown in the following table.

1. You must choose from these groups to determine which ones will be configured in static-12x mode.
2. Record the lowest numbered port in each group. For example, for group 3, you would record 7 and for group 7, you would record 19.

Table 7. 7048-120 or SFS7000P

Group 1			**Group 2**			**Group 3**			**Group 4**		
1	2	3	4	5	6	7	8	9	10	11	12
Group 5			**Group 6**			**Group 7**			**Group 8**		
13	14	15	16	17	18	19	20	21	22	23	24

Determining groups to use for static-12x mode, 7048-270 or SFS7008P switch

For a 7048-270 or SFS7008P, each 4x Line Interface Module (LIM) has 12 ports. These ports are logically divided into groups of three on a modulo-3 boundary, as shown in the following table.

1. You must determine which LIM(s) and which group(s) on the chosen LIM(s) you want to be configured for static-12x mode.
2. Record the LIM number.
3. Record the lowest numbered port in each group. For example, for group 2, you would record 4, and for group 4 you would record 10.

Table 8. 7048-270 or SFS7008P

Group 1			**Group 2**			**Group 3**			**Group 4**		
1	2	3	4	5	6	7	8	9	10	11	12

Configuring Static-12x groups

To configure a group to static-12x mode, you only have to set the lowest numbered port to an admin speed of 12x with auto-negotiate disabled. Regardless of which port will become Active in the static-12x mode, the lowest numbered port is always the one that determines if the group is in static-12x mode.

To configure the static-12x groups, you must do so through the switch command line interface (CLI). Connecting to and logging onto the switch CLI is described in the Switch Install Instructions, See "Accessing an InfiniBand switch command line interface" on page 158 and the switch hardware guides.

The cable connection information planned here should be recorded in a “Switch planning worksheet” on page 28 for switch port connections and in a Server planning worksheet for HCA port connections.

Planning octopus cables in static 12x cabling ends here.

Planning 4x connections for octopus cables

Plan to connect the three 4x octopus cable ends first, and then connect the 12x cable end.

Before you begin, you should have determined the groups that you want to use and you should have configured those groups to static-12x mode.

You are now ready to connect the 4x connectors of the Octopus cables to the switch ports. It is a best practice to do this before connecting the 12x connector of the cable to the 12x device.

Important: The ordering of these connections is critical in static-12x mode. You must connect the 4x connectors labeled 16, 17 and 18 to the correct ports in each group. Otherwise, the links will not be able to pass data.

See the following 7048-120, SFS7000P, 7048-270, and SFS7008P cabling tables.

7048-120 or SFS7000P cabling tables

The following table lists the connections for each group in a 7048-120 or SFS7000P.

1. Find the port numbers for the port group listed in the table.
2. Using the order shown in the table, connect the corresponding 4x octopus cable end connectors to the appropriate ports on the switch.

For example, if you are cabling ports 7-9, port 7 receives connector 16, port 8 receives connector 17, and port 9 receive connector 18.

Table 9. 7048-120 or SFS7000P cable ordering for octopus cable connectors

Switch port row	Port group numbers	4x octopus cable end connector number
Top ports of switch	Ports 1-3	16, 17, 18
	Ports 4-6	16, 17, 18
	Ports 7-9	16, 17, 18
	Ports 10-12	16, 17, 18
Bottom Ports of Switch	Ports 13-15	18, 17, 16
	Ports 16-18	18, 17, 16
	Ports 19-21	18, 17, 16
	Ports 22-24	18, 17, 16

7048-270 or SFS7008P cabling tables

The following tables lists the connections for each group in a 7048-270 orSFS7008P.

1. Find the port numbers for the port group and LIM listed in the table.
2. Using the order shown in the table, connect the corresponding 4x octopus cable end connectors to the appropriate ports on the switch.

For example, if you are cabling ports 7-9 on LIM 2, port 7 receives connector 18, port 8 receives connector 17, and port 9 receive connector 16. Or, if you’re cabling the ports 4-6 on LIM 6, port 4 receives 16, port 5

receives 17 and port 6 receives 18.

Table 10. 7048-270 or SFS7008P cable ordering for octopus cable connectors

LIM	Octopus cable connectors for group ports 1-3	Octopus cable connectors for group ports 4-6	Octopus cable connectors for group ports 7-9	Octopus cable connectors for group ports 10-12
1	18, 17, 16	18, 17, 16	16, 17, 18	16, 17, 18
2	16, 17, 18	16, 17, 18	18, 17, 16	18, 17, 16
3	16, 17, 18	16, 17, 18	18, 17, 16	18, 17, 16
4	18, 17, 16	18, 17, 16	16, 17, 18	16, 17, 18
5	18, 17, 16	18, 17, 16	16, 17, 18	16, 17, 18
6	16, 17, 18	16, 17, 18	18, 17, 16	18, 17, 16
7	16, 17, 18	16, 17, 18	18, 17, 16	18, 17, 16
8	18, 17, 16	18, 17, 16	16, 17, 18	16, 17, 18

Note: LIMs 1, 4, 5, and 8 have the same connector pattern. LIMs 2, 3, 6, and 7 have the same connector pattern.

Planning 4x connections for octopus cables ends here.

Planning Aids

The following lists some optional tasks that can help you when planning your cluster hardware.

- Determine a convention for frame numbering and slot numbering, where slots are the locations of units as you go from the bottom of the frame to the top. If you have empty space in a frame, reserve a number for that space.
- Determine a convention for switch and system unit naming that includes the physical location of units including their frame numbers and slot numberss.
- Prepare labels for frames to indicate frame numbers.
- Prepare cable labels for each end of the cables. Indicate the ports to which each end of the cable is to connect.
- Document where switches and servers are located and which HMCs manage them.
- Print a floor plan and keep it with the HMCs.

Planning Aids ends here.

Planning checklist

The planning checklist helps you track your progress through the planning process.

All the worksheets and checklists are available to be copied in Appendix B. Planning and Installation Worksheets, beginning with “Installation coordinating worksheets” on page 241.

Table 11. Planning checklist

Step	Target	Complete
Start planning checklist		
Gather documentation and review planning information for individual devices.		

Table 11. Planning checklist (continued)

Step	Target	Complete
Ensure that you have planned for: • Servers • I/O devices • InfiniBand network devices • Frames or racks • Service network, including: – HMCs – Ethernet devices – CSM Management Server (for multiple HMC environments) – AIX NIM server (for servers with no removable media) – Linux distribution server (for servers with no removable media) • System management applications (HMC and CSM and IBM Network Manager) • Physical dimension and weight characteristics • Electrical characteristics • Cooling characteristics		
Ensure that you have the "Required levels of support, firmware, and devices" on page 7 for your network hardware		
Review cabling and topology documentation for InfiniBand networks		
Review "Installation flow of clusters using the IBM Network Manager" on page 11		
Review "Planning for a high-performance computing message-passing interface configuration" on page 18		
Review "Planning octopus cables in static 12x cabling" on page 19		
Review "Planning Aids" on page 22		
Complete planning worksheets		
Complete planning process		
Review readme files and online information related to the software and firmware to ensure that you have up-to-date information and the latest supported levels		

Planning checklist ends here.

Planning worksheets

The planning worksheets are used when you are planning your cluster.

Tip: It is best to keep the sheets somewhere that is accessible to the system administrators and service representatives not only during the installation process, but also for future reference during maintenance, upgrade, or repair actions.

All the worksheets and checklists are available to be copied in Appendix B. Planning and Installation Worksheets, beginning with "Installation coordinating worksheets" on page 241.

Using planning worksheets

The worksheets do not cover all situations (especially with regard to the number of instances of slots in a frame, servers in a frame, or I/O slots in a server). However, they provide enough of a base upon which you can build a custom worksheet for your application. In some cases, you might find it useful to bring the worksheets into a spreadsheet so that you may fill out repetitive information. Otherwise, you can

devise a method to indicate repetitive information in a formula on printed worksheets so that you do not have to complete large numbers of worksheets for a large cluster that is likely to have a definite pattern in frame, server, and switch configuration.

To access the worksheets, "Cluster summary worksheet" on page 25. Complete the worksheets in the following sequence:

1. "Cluster summary worksheet" on page 25
2. "Frame and rack planning worksheet" on page 26
3. "Server planning worksheet" on page 27
4. "Switch planning worksheet" on page 28

For an example of how to complete the worksheets, see "Worksheet examples" on page 32.

All the worksheets are available to be copied in Appendix B. Planning and Installation Worksheets, beginning with "Installation coordinating worksheets" on page 241.

Cluster summary worksheet

Record information about the items that are in your network.

Cluster summary worksheet
Cluster name:
Application:
Number and types of servers:
Number of servers and HCAs per server: **Note:** If there are servers with various numbers of HCAs, list the number of servers with each configuration; for example, 12 servers with one 2-port HCA; 4 servers with two 2-port HCAs.
Number of 7048-120 or SFS7000P switches:
Number of 7048-270 or SFS7008P switches:
Number of subnets:
List of GID-prefixes and subnet masters (assign a number to a subnet for easy reference):
Switch partitions:
Number of frames:
Number of HMCs:
CSM and Cluster Ready Hardware Server used?
Number of Service Ethernet networks:
Service network domains:
Service network DHCP server locations:
Service network switches with static IP:
Service network HMCs with static IP:
Service network DHCP range(s):
AIX NIM server info:
Linux distribution server info:
Power requirements:
Maximum cooling required:
Number of cooling zones:
Maximum weight per square foot:

Frame and rack planning worksheet

The frame and rack planning worksheet is used for planning how to populate frames or racks. You will likely have to understand the numbers of each device type (servers, switches, and bulk power adapters). For the slots, indicate the range of slots or drawers that a device will populate. A standard method for naming slots can either be found in the documentation for the frames or servers, or use EIA heights (1.75 inches) as a standard.

Frame and rack planning worksheet		
Frame or rack number(s): _____ Frame MTM or feature: _____ Frame size: _____ (19-inch or 24-inch) Number of slots: _____		
Slots	**Device type (server, switch, or BPA,) Indicate Machine type and model**	**Device name**

Server planning worksheet

You can use this worksheet as a template for multiple servers with similar configurations. For such cases, you will want to give the range of names of these servers and where they are located. Also, you can use the configuration note to remind you of other specific characteristics of the server. It is important to note the type of HCAs to be used.

Server planning worksheet				
Name(s): _____ Type(s): _____ Frame(s)/slot(s): _____ Number and type of HCAs_____ Num LPARs/LHCAs: _____ IP addressing: _____ IP addressing of service subsystem: _____ LPAR IP addressing: _____ MPI addressing: _____ Configuration notes:				
HCA information				
HCA	Capability (Sharing)	HCA port	Switch connection	GID prefix
LPAR information				
LPAR/LHCA (give name)	OS Type	GUID index	Shared HCA (capability)	Switch partition

Switch planning worksheet

There is a worksheet in this section for each type of switch.

When documenting connections to switch ports, it is suggested that you note both a shorthand for your own use and the IBM locations.

For example, if you are connecting port 1 of a 7048-120 switch to port 1 of the only HCA in a 9118-575 server, that you are going to name f1n1, you might use the shorthand f1n1-HCA1-Port1 to indicate this connection.

It would be useful to also note the IBM location code for this HCA port, as well. You can get the location code information specific to each server in the server's documentation and do this at the time of planning, or you can work with the IBM Service Representative at the time of the installation to make the proper notation with regard to the IBM location code. Generally, the only piece of information that is not available during the planning phase is the server's serial number, which is used as part of the location code.

HCAs generally have the location code: U[*server or feature code*].001.[*server's serial number*]-P*x*-C*y*-T*z*; where:
P*x* represents the planar into which the HCA plugs.
C*y* represents the planar connector into which the HCA plugs.
T*z* represents the HCA's port into which the cable plugs.

Note: The groupings for octopus cables used on 4x switches connected to 12x devices (such as HCAs) are indicated by a bold line around the associated ports. You will also note that one port is indicated as the port to which to connect the 4x cable connector labelled 16. On the switch worksheet for these ports you will see (16) after the applicable port number. In the following example of the first group of three ports, notice the (16) which indicates that octopus cable connector 16 would connect to port 1.

Ports	**Connection**
1(16)	
2	
3	

Use the following worksheet for planning 7048-270 and SFS7008P switches.

7048-270 and SFS7008P switch worksheet	
Switch MTM: _____ (7048-270 or SFS7008P) Switch name: _____ Frame and slot: _____ IP address: _____ (if single HMC in cluster, indicate DHCP) GID-prefix: _____ LMC: _____ (0=default; 2=if used in HPC cluster) Switch MTMS: _____ (Complete during installation)	
Ports	**Connection**
1 (16)	
2	
3	
4 (16)	
5	
6	
7 (16)	
8	
9	
10 (16)	
11	
12	
13	
14	
15 (16)	
16	
17	
18 (16)	
19	
20	
21 (16)	
22	
23	
24 (16)	

Use the following worksheet for planning 7048-270 and SFS7008P switches.

7048-270 and SFS7008P switch worksheet			
Switch MTM: _____ (7048-270 or SFS7008P) Switch name: _____ Frame and slot: _____ IP address: _____ (if single HMC in cluster, indicate DHCP) GID-prefix: _____ LMC: _____ (0=default; 2=if used in HPC cluster) Switch MTMS: _____ (Complete during installation)			
LIM 1		**LIM 2**	
Ports	**Connection**	**Ports**	**Connection**
1		1 (16)	
2		2	
3 (16)		3	
4		4 (16)	
5		5	
6 (16)		6	
7 (16)		7	
8		8	
9		9 (16)	
10 (16)		10	
11		11	
12		12 (16)	
LIM 3		**LIM 4**	
Ports	**Connection**	**Ports**	**Connection**
1 (16)		1	
2		2	
3		3 (16)	
4 (16)		4	
5		5	
6		6 (16)	
7		7 (16)	
8		8	
9 (16)		9	
10		10 (16)	
11		11	
12 (16)		12	
LIM 5		**LIM 6**	
Ports	**Connection**	**Ports**	**Connection**
1		1 (16)	
2		2	
3 (16)		3	
4		4 (16)	

7048-270 and SFS7008P switch worksheet			
5		5	
6 (16)		6	
7 (16)		7	
8		8	
9		9 (16)	
10 (16)		10	
11		11	
12		12 (16)	
LIM 7		**LIM 8**	
Ports	**Connection**	**Ports**	**Connection**
1 (16)		1	
2		2	
3		3 (16)	
4 (16)		4	
5		5	
6		6 (16)	
7		7 (16)	
8		8	
9 (16)		9	
10		10 (16)	
11		11	
12 (16)		12	

Worksheet examples

The following examples are completed worksheets for a 96 way cluster with two subnets. There are several types of servers used.

Cluster summary worksheet
Cluster name: Example (ez)
Application: HPC
Number of servers and HCAs per server: (94) 9118-575 (2) 9119-590s Number of HCAs/server: (96) with two 2-port HCA (using extra HCAs for future expansion) **Note:** If there are servers with various numbers of HCAs, list the number of servers with each configuration; for example, 12 servers with one 2-port HCA; 4 servers with two 2-port HCAs.
Number of 7048-120 or SFS7000P switches: 0
Number of 7048-270 or SFS7008P switches: 2
Number of subnets: 2
List of GID-prefixes and subnet masters (assign a number to a subnet for easy reference): subnet1 = FE:80:00:00:00:00:00:00 (ezfr11sw1) subnet 2 =FE:80:00:00:00:00:00:01 (ezfr11sw2)
Switch partitions:
Number of frames: _(8) 575 24″; (2) 590 24″
Number of HMCs: 3
CSM and Cluster Ready Hardware Server used? yes
Number of Ethernet Service networks: 2
Service network domains: 10.0.1 and 10.0.2
Service network DHCP server locations: CSM MS
Service network switches with static IP: 10.0.2.8 and 10.0.2.9
Service network HMCs with static IP: 10.0.[2,3].2, 10.0.[2,3].3, 10.0.[2,3].4
Service network DHCP range(s): 10.0.2.[10-106] and 10.0.3.[10-106]
AIX NIM server info: nim-server1.[mydomain]
Linux distribution server info: linux-dist1.[mydomain]
Power requirements: (8) 41.6 kW @380V (9118-575) + (2) 22.7kW (9119-590) @ 380V (2) 600W @120V (SFS7008P)
Maximum cooling required: 3000 ft3 @ 62F (9118-575) _2000 ft3@62F (9119-590 w/ 4 I/O drawers)
Number of cooling zones: 1
Maximum weight: 3479lbs/16.44ft2 (9118-575)

Frame planning worksheet:

The following completed worksheet example is for frame planning.

Frame planning worksheet (1/3)		
Frame number(s): _1-7_____ Frame MTM/feature: __FC5793_____ Frame size: ____24-inch_____ (19-inch or 24-inch) Number of slots: _____12_____		
Slots	Device Type (server, switch, BPA, etc...) Indicate MTM	Device Name
1	Server; 9118-575	ezfr[x]n1; where x is the frame number 1-7 (ez = example cluster name)
2	Server; 9118-575	ezfr[x]n2; where x is the frame number 1-7
3	Server; 9118-575	ezfr[x]n3; where x is the frame number 1-7
4	Server; 9118-575	ezfr[x]n4; where x is the frame number 1-7
5	Server; 9118-575	ezfr[x]n5; where x is the frame number 1-7
6	Server; 9118-575	ezfr[x]n6; where x is the frame number 1-7
7	Server; 9118-575	ezfr[x]n7; where x is the frame number 1-7
8	Server; 9118-575	ezfr[x]n8; where x is the frame number 1-7
9	Server; 9118-575	ezfr[x]n9; where x is the frame number 1-7
10	Server; 9118-575	ezfr[x]n10; where x is the frame number 1-7
11	Server; 9118-575	ezfr[x]n11; where x is the frame number 1-7
12	Server; 9118-575	ezfr[x]n12; where x is the frame number 1-7

Frame planning worksheet (2/3)		
Frame number(s): _8-9_____ Frame MTM/feature: __9119-590 base rack_____ Frame size: ____24-inch_____ (19-inch or 24-inch) Number of slots: _____1_____		
Slots	**Device type** (server, switch, BPA, etc...) Indicate MTM	**Device name**
1	Server; 9119-590	ezfr[x]n1; where x is the frame number (8 or 9) (ez = example cluster name)

Frame planning worksheet (3/3)		
Frame number(s): _11_____ Frame MTM/feature: __7014-T00 _____ Frame size: ____19-inch_____ (19-inch or 24-inch) Number of slots: _____2_____		
Slots	**Device type** (server, switch, BPA, etc...) Indicate MTM	**Device name**
1-2	Switch; SFS7008P	ezfr11sw[x]; where x is the slot number (1 or 2) (ez = example cluster name)

Server planning:

The following completed worksheet example is for server planning.

Server planning worksheet (1/1)				
Name(s): _ezfr[1-8]n[1-12]__(except ezfr8n[11-12] do not exist)_____ Type(s): ___9118-575_____ Frame(s)/slot(s): ____frame 1-8; slot 11-12 (except frame 8 slots 11 & 12 are not populated)___ Number and type of HCAs_____2 GX bus (no PCI)_____ Num LPARs/LHCAs: _____2 per node_____ IP addressing: __10.0.0.[10-106] _____ IP addressing of service subsystem: ___DHCP from CSM MS (10.0.2.[10-106] and 10.0.3.[10-106])__ LPAR IP addressing: __9.117.25.[100-196]_____ MPI addressing: _____10.0.1.[110-206]_____ Configuration Note:				
HCA information				
HCA	Capability (sharing) level	HCA port	Switch connection	GID prefix
1	Low	1	ezfr11sw1; start with ezfr1n1 through to ezfr10n1 from LIM 1 port 1 through LIM8 port 12.	FE:80:00:00:00:00:00:00
		2	No connect	
2	Low	1	ezfr11sw2; start with ezfr1n1 through to ezfr10n1 from LIM 1 port 1 through LIM8 port 12.	FE:80:00:00:00:00:00:01
		2	No connect	
LPAR information				
LPAR/LHCA (give name)	OS type	GUID index	Shared HCA	Switch partition
ezfrXnYsq01	Linux	1	1	N/A
ezfrXnYsq02	Linux	2	2	N/A

7048-270 and SFS7008P switch worksheet (1/2)

Switch MTM: ____SFS7008P_____ (7048-270 or SFS7008P)
Switch name: ____ezfr11sw1_____
Frame and slot: ___frame 11; slot 1_____
IP address: _____10.0.2.8_____ (if single HMC in cluster, indicate DHCP)
GID-prefix: _____FE:80:00:00:00:00:00:00_____
LMC: _____2_____ (0=default; 2=if used in HPC cluster)
Switch MTMS: ___7048-708-2432751_____ (Complete during installation)

LIM 1		**LIM 2**	
Ports	**Connection**	**Ports**	**Connection**
1	ezfr1n1-HCA1-Port1 U787C.001.1234567-P2-C66-T1	1 (16)	ezfr2n1-HCA1-Port1
2	ezfr1n2-HCA1-Port1 U787C.001.1234568-P2-C66-T1	2	ezfr2n2-HCA1-Port1
3 (16)	ezfr1n3-HCA1-Port1 U787C.001.1234569-P2-C66-T1	3	ezfr2n3-HCA1-Port1
4	ezfr1n4-HCA1-Port1 U787C.001.1234560-P2-C66-T1	4 (16)	ezfr2n4-HCA1-Port1
5	ezfr1n5-HCA1-Port1	5	ezfr2n5-HCA1-Port1
6 (16)	ezfr1n6-HCA1-Port1	6	ezfr2n6-HCA1-Port1
7 (16)	ezfr1n7-HCA1-Port1	7	ezfr2n7-HCA1-Port1
8	ezfr1n8-HCA1-Port1	8	ezfr2n8-HCA1-Port1
9	ezfr1n9-HCA1-Port1	9 (16)	ezfr2n9-HCA1-Port1
10 (16)	ezfr1n10-HCA1-Port1	10	ezfr2n10-HCA1-Port1
11	ezfr1n11-HCA1-Port1	11	ezfr2n11-HCA1-Port1
12	ezfr1n12-HCA1-Port1	12 (16)	ezfr2n12-HCA1-Port1
LIM 3		**LIM 4**	
Ports	**Connection**	**Ports**	**Connection**
1 (16)	ezfr3n1-HCA1-Port1	1	ezfr4n1-HCA1-Port1
2	ezfr3n2-HCA1-Port1	2	ezfr4n2-HCA1-Port1
3	ezfr3n3-HCA1-Port1	3 (16)	ezfr4n3-HCA1-Port1
4 (16)	ezfr3n4-HCA1-Port1	4	ezfr4n4-HCA1-Port1
5	ezfr3n5-HCA1-Port1	5	ezfr4n5-HCA1-Port1
6	ezfr3n6-HCA1-Port1	6 (16)	ezfr4n6-HCA1-Port1
7	ezfr3n7-HCA1-Port1	7 (16)	ezfr4n7-HCA1-Port1
8	ezfr3n8-HCA1-Port1	8	ezfr4n8-HCA1-Port1
9 (16)	ezfr3n9-HCA1-Port1	9	ezfr4n9-HCA1-Port1
10	ezfr3n10-HCA1-Port1	10 (16)	ezfr4n10-HCA1-Port1
11	ezfr3n11-HCA1-Port1	11	ezfr4n11-HCA1-Port1
12 (16)	ezfr3n12-HCA1-Port1	12	ezfr4n12-HCA1-Port1
LIM 5		**LIM 6**	
Ports	**Connection**	**Ports**	**Connection**
1	ezfr5n1-HCA1-Port1	1 (16)	ezfr6n1-HCA1-Port1
2	ezfr5n2-HCA1-Port1	2	ezfr6n2-HCA1-Port1
3 (16)	ezfr5n3-HCA1-Port1	3	ezfr6n3-HCA1-Port1

7048-270 and SFS7008P switch worksheet (1/2)			
4	ezfr5n4-HCA1-Port1	4 (16)	ezfr6n4-HCA1-Port1
5	ezfr5n5-HCA1-Port1	5	ezfr6n5-HCA1-Port1
6 (16)	ezfr5n6-HCA1-Port1	6	ezfr6n6-HCA1-Port1
7 (16)	ezfr5n7-HCA1-Port1	7	ezfr6n7-HCA1-Port1
8	ezfr5n8-HCA1-Port1	8	ezfr6n8-HCA1-Port1
9	ezfr5n9-HCA1-Port1	9 (16)	ezfr6n9-HCA1-Port1
10 (16)	ezfr5n10-HCA1-Port1	10	ezfr6n10-HCA1-Port1
11	ezfr5n11-HCA1-Port1	11	ezfr6n11-HCA1-Port1
12	ezfr5n12-HCA1-Port1	12 (16)	ezfr6n12-HCA1-Port1
LIM 7		**LIM 8**	
Ports	**Connection**	**Ports**	**Connection**
1 (16)	ezfr7n1-HCA1-Port1	1	ezfr8n1-HCA1-Port1
2	ezfr7n2-HCA1-Port1	2	ezfr8n2-HCA1-Port1
3	ezfr7n3-HCA1-Port1	3 (16)	ezfr8n3-HCA1-Port1
4 (16)	ezfr7n4-HCA1-Port1	4	ezfr8n4-HCA1-Port1
5	ezfr7n5-HCA1-Port1	5	ezfr8n5-HCA1-Port1
6	ezfr7n6-HCA1-Port1	6 (16)	ezfr8n6-HCA1-Port1
7	ezfr7n7-HCA1-Port1	7 (16)	ezfr8n7-HCA1-Port1
8	ezfr7n8-HCA1-Port1	8	ezfr8n8-HCA1-Port1
9 (16)	ezfr7n9-HCA1-Port1	9	ezfr8n9-HCA1-Port1
10	ezfr7n10-HCA1-Port1	10 (16)	ezfr8n10-HCA1-Port1
11	ezfr7n11-HCA1-Port1	11	ezfr9n1-HCA1-Port1 U787C.001.6789321-P2-C8-T1
12 (16)	ezfr7n12-HCA1-Port1	12	ezfr10n1-HCA1-Port1 U787C.001.6789321-P2-C8-T1

7048-270 and SFS7008P switch worksheet (2/2)

Switch MTM: ____SFS7008P_____ (7048-270 or SFS7008P)
Switch name: ____ezfr11sw2_____
Frame and slot: ___frame 11; slot 2_____
IP address: _____10.0.2.9_____ (if single HMC in cluster, indicate DHCP)
GID-prefix: _____FE:80:00:00:00:00:00:01_____
LMC: _____2_____ (0=default; 2=if used in HPC cluster)
Switch MTMS: ___SFS7008P-2432751_____ (Complete during installation)

LIM 1		**LIM 2**	
Ports	**Connection**	**Ports**	**Connection**
1	ezfr1n1-HCA2-Port1 U787C.001.1234567-P2-C65-T1	1 (16)	ezfr2n1-HCA2-Port2
2	ezfr1n2-HCA1-Port1 U787C.001.1234568-P2-C65-T1	2	ezfr2n2-HCA2-Port2
3 (16)	ezfr1n3-HCA2-Port1 U787C.001.1234569-P2-C65-T1	3	ezfr2n3-HCA2-Port2
4	ezfr1n4-HCA2-Port1 U787C.001.1234560-P2-C65-T1	4 (16)	ezfr2n4-HCA2-Port2
5	ezfr1n5-HCA2-Port2	5	ezfr2n5-HCA2-Port2
6 (16)	ezfr1n6-HCA2-Port2	6	ezfr2n6-HCA2-Port2
7 (16)	ezfr1n7-HCA2-Port2	7	ezfr2n7-HCA2-Port2
8	ezfr1n8-HCA2-Port2	8	ezfr2n8-HCA2-Port2
9	ezfr1n9-HCA2-Port2	9 (16)	ezfr2n9-HCA2-Port2
10 (16)	ezfr1n10-HCA2-Port2	10	ezfr2n10-HCA2-Port2
11	ezfr1n11-HCA2-Port2	11	ezfr2n11-HCA2-Port2
12	ezfr1n12-HCA2-Port2	12 (16)	ezfr2n12-HCA2-Port2
LIM 3		**LIM 4**	
Ports	**Connection**	**Ports**	**Connection**
1 (16)	ezfr3n1-HCA2-Port2	1	ezfr4n1-HCA2-Port2
2	ezfr3n2-HCA2-Port2	2	ezfr4n2-HCA2-Port2
3	ezfr3n3-HCA2-Port2	3 (16)	ezfr4n3-HCA2-Port2
4 (16)	ezfr3n4-HCA2-Port2	4	ezfr4n4-HCA2-Port2
5	ezfr3n5-HCA2-Port2	5	ezfr4n5-HCA2-Port2
6	ezfr3n6-HCA2-Port2	6 (16)	ezfr4n6-HCA2-Port2
7	ezfr3n7-HCA2-Port2	7 (16)	ezfr4n7-HCA2-Port2
8	ezfr3n8-HCA2-Port2	8	ezfr4n8-HCA2-Port2
9 (16)	ezfr3n9-HCA2-Port2	9	ezfr4n9-HCA2-Port2
10	ezfr3n10-HCA2-Port2	10 (16)	ezfr4n10-HCA2-Port2
11	ezfr3n11-HCA2-Port2	11	ezfr4n11-HCA2-Port2
12 (16)	ezfr3n12-HCA2-Port2	12	ezfr4n12-HCA2-Port2
LIM 5		**LIM 6**	
Ports	**Connection**	**Ports**	**Connection**
1	ezfr5n1-HCA2-Port2	1 (16)	ezfr6n1-HCA2-Port2
2	ezfr5n2-HCA2-Port2	2	ezfr6n2-HCA2-Port2
3 (16)	ezfr5n3-HCA2-Port2	3	ezfr6n3-HCA2-Port2

7048-270 and SFS7008P switch worksheet (2/2)			
4	ezfr5n4-HCA2-Port2	4 (16)	ezfr6n4-HCA2-Port2
5	ezfr5n5-HCA2-Port2	5	ezfr6n5-HCA2-Port2
6 (16)	ezfr5n6-HCA2-Port2	6	ezfr6n6-HCA2-Port2
7 (16)	ezfr5n7-HCA2-Port2	7	ezfr6n7-HCA2-Port2
8	ezfr5n8-HCA2-Port2	8	ezfr6n8-HCA2-Port2
9	ezfr5n9-HCA2-Port2	9 (16)	ezfr6n9-HCA2-Port2
10 (16)	ezfr5n10-HCA2-Port2	10	ezfr6n10-HCA2-Port2
11	ezfr5n11-HCA2-Port2	11	ezfr6n11-HCA2-Port2
12	ezfr5n12-HCA2-Port2	12 (16)	ezfr6n12-HCA2-Port2
LIM 7		**LIM 8**	
Ports	**Connection**	**Ports**	**Connection**
1 (16)	ezfr7n1-HCA2-Port2	1	ezfr8n1-HCA2-Port2
2	ezfr7n2-HCA2-Port2	2	ezfr8n2-HCA2-Port2
3	ezfr7n3-HCA2-Port2	3 (16)	ezfr8n3-HCA2-Port2
4 (16)	ezfr7n4-HCA2-Port2	4	ezfr8n4-HCA2-Port2
5	ezfr7n5-HCA2-Port2	5	ezfr8n5-HCA2-Port2
6	ezfr7n6-HCA2-Port2	6 (16)	ezfr8n6-HCA2-Port2
7	ezfr7n7-HCA2-Port2	7 (16)	ezfr8n7-HCA2-Port2
8	ezfr7n8-HCA2-Port2	8	ezfr8n8-HCA2-Port2
9 (16)	ezfr7n9-HCA2-Port2	9	ezfr8n9-HCA2-Port2
10	ezfr7n10-HCA2-Port2	10 (16)	ezfr8n10-HCA2-Port2
11	ezfr7n11-HCA2-Port2	11	ezfr9n1-HCA2-Port2 U787C.001.6789321-P2-C8-T1
12 (16)	ezfr7n12-HCA2-Port2	12	ezfr10n1-HCA2-Port2 U787C.001.6789321-P2-C8-T1

Installing a cluster that has an InfiniBand network

After planning the cluster and InfiniBand network, when your hardware is available, use the following procedures to install your InfiniBand cluster.

This documentation does not cover installation of I/O devices other than those in the InfiniBand network. All non-InfiniBandI/O devices are considered part of the server installation procedure.

1. Separate the installation tasks based on the generalized tasks and the people responsible for them, as outlined in "Installation responsibilities by organization" on page 12 and "Installation responsibilities by units and devices" on page 13.
2. Ensure that you understand the "Installation flow of clusters using the IBM Network Manager" on page 11 in the "Planning for InfiniBand networks" on page 5 section. Pay close attention to the merge points that are crucial to the coordination of a successful installation.
3. To complete the installation of your InfiniBand network, use the detailed processes that follow. Perform the following tasks in the order shown:

   Note: Use the following links to find detailed procedures for each task that is necessary to install the network.

   Perform each task in the order shown, then go to the next task. The major task numbers found in the "Order of Installation" on page 13 are referenced in the detailed procedures.

   a. "Setup the site for power, cooling, and floor requirements" on page 41
   b. "Install and configure the management subsystem" on page 41
   c. "Installing and configuring the cluster servers" on page 47
   d. "InfiniBand switch installation and configuration for vendor switches" on page 55
   e. "Attach cables to the InfiniBand network" on page 59
   f. "Verify the InfiniBand network topology and operation" on page 62
4. If you are expanding your network or cluster, some major tasks may not apply at all, some may apply in their entirety, and others may only be partially required. First, use the following table to determine which major tasks apply. If there are deviations based on an expansion scenario, the individual major task will indicate the deviations.

Table 12.

Detailed procedure	Adding InfiniBand hardware to an existing cluster (switches and HCAs)	Adding new servers to an existing InfiniBand network	Adding HCAs to an existing InfiniBand network	Adding a subnet to an existing InfiniBand network	Adding servers and a subnet to an existing InfiniBand network
"Setup the site for power, cooling, and floor requirements" on page 41	Yes	Yes	floor tile cutouts for cables	Yes	Yes
"Install and configure the management subsystem" on page 41	Yes	Yes (for install images)	No	Yes	Yes

© Copyright IBM Corp. 2004, 2007

Table 12. (continued)

Detailed procedure	Adding InfiniBand hardware to an existing cluster (switches and HCAs)	Adding new servers to an existing InfiniBand network	Adding HCAs to an existing InfiniBand network	Adding a subnet to an existing InfiniBand network	Adding servers and a subnet to an existing InfiniBand network
"Installing and configuring the cluster servers" on page 47	Yes	If CSM and CRHS* must be added**	No	Yes	Yes
"InfiniBand switch installation and configuration for vendor switches" on page 55	Yes	Yes	Yes	Yes	Yes
"Attach cables to the InfiniBand network" on page 59	Yes	Yes	Yes	Yes	Yes
"Verify the InfiniBand network topology and operation" on page 62	Yes	Yes	Yes	Yes	Yes

* Cluster-ready hardware server (CRHS)
** This occurs when:

- A single HMC is in an existing cluster
- Existing switches are configured to use DHCP service on the service Ethernet network
- Servers are being added to an existing cluster
- Added servers require you to add one or more new HMCs
- You must use CSM and CRHS, and configure the switches with static IP addressing on the service Ethernet network.

Each of the above major task sections contains information to describe the following:

- An overview of the task
- An overview of expansion scenarios versus new cluster installation
- Installation responsibilities section
- Reference documentation section
- Installation tasks section
  - If the installation procedure might be performed by a non-IBM vendor, then there will be a list of key points that cover important considerations for configuration and order of installation with which the non-IBM vendor must comply for an installation in a cluster. This allows the non-IBM vendor to use alternative procedures as long as they meet the key points.
  - In the detailed procedure, task reference numbers cross-reference the detailed procedure steps to the Figure 3 on page 15 figure. Each procedure contains its list of major task references from the high-level cluster installation flow figure. Furthermore, a specific major task reference is listed at the beginning of major steps in the procedure.

Pay close attention to any pre-requisites listed for the procedure or steps in the procedure

IBM Service representative installation responsibilities

IBM Service installation responsibilities include installing IBM machine types that are IBM installable versus those that are customer installable. In addition to the typical repair responsibilities during installation, IBM service is responsible for repairing the InfiniBand cables and HCAs.

IBM Service is not responsible for installing the InfiniBand switches, nor the InfiniBand cables.

Setup the site for power, cooling, and floor requirements

The site setup for power, cooling and the floor requirements encompasses major task **U1** illustrated in Figure 3 on page 15.

The setup for the power, cooling and floor requirements must be complete before proceeding to install the cluster. The setup must meet all documented requirements for the individual units, frames, systems and adapters in the cluster. Generally this is performed by the customer, IBM installation planning representative, or a contractor.

Note: If you are installing HCAs into existing servers, you should only have to perform operations involving cable routing and floor tile cut-outs.

Install and configure the management subsystem

The management subsystem installation and configuration encompasses major tasks **M1** - **M4** , which are illustrated in Figure 3 on page 15.

You will be installing and configuring HMCs, an Ethernet service network, and possibly a CSM management server, as well as building an AIX NIM SPoT to run eServer diagnostics for servers without removable media (CD and DVD drives). The eServer diagnostics are available only in AIX, and you will require an AIX NIM SPoT even if your partitions are running another operating system such as Linux. If you plan to install servers without removable media, you will also need a Linux distribution server for distribution of the operating system.

This is not a detailed description of how to install the management subsystem components, because such procedures are described in detail in documentation for the individual components, such as devices and applications. This topic collection describes the order of installation and key points that you need to consider when installing and configuring the management consoles.

The management consoles that are to be installed in this topic collection are the HMC and possibly a CSM management console. Because the management consoles are the heart of the service network, they are key to successfully installing and configuring the cluster. Before you do any significant bring-up and configuration, these devices must be installed and configured so that they are ready to discover, connect, and manage the rest of the devices in the cluster.

Important: If possible, complete this procedure (through the final configuration step) before beginning other procedures, such as installing cluster servers or switches. This will alleviate the situation where various installation personnel may be waiting on site for key parts of this procedure to be completed. Depending on the arrival of units on site, this is not always practical. Therefore, it is important to review the "Order of Installation" on page 13 and Figure 3 on page 15 figure from "Planning for InfiniBand networks" on page 5 networks to identify the merge points where a step in a major task or procedure being performed by one person is dependent on the completion of steps in another major task or procedure being performed by another person.

Expanding and configuring the management subsystem

If you are expanding InfiniBand network capabilities for an existing cluster, then you might need

to approach the management subsystem installation and configuration differently than with a new cluster installation. The process for the management subsystem installation and configuration task is based on a new cluster installation, but it will indicate where there are variances for expansion scenarios. The following table shows how the new cluster installation process is changed when you are working in an expansion scenario:

Table 13. Expansion of an existing cluster and how it differs from a new installation

Scenario	Effects
Adding InfiniBand hardware (switches and HCAs) to an existing cluster	• Connect the InfiniBand switch to the Ethernet service network • Might require additional service network Ethernet switches or routers to accommodate new InfiniBand switches • If there are multiple HMCs in the existing cluster, you must have CSM and cluster-ready hardware server.
Adding new servers to an existing InfiniBand switch network	• Connect the new server to the Ethernet service network • Ensure operating system distribution capability for new servers without removable media • Might require additional HMCs to accommodate the new servers. If expanding beyond a single HMC or adding a CSM management server and cluster ready hardware server, you will need to deconfigure the current IP addressing method and reconfigure the network using DHCP on the CSM management server. • Might require additional Ethernet service network switches or routers to accommodate the new servers
Adding HCAs to an existing InfiniBand network	This should not affect the Ethernet service network.
Adding a subnet to an existing InfiniBand network	• Connect the InfiniBand switch to the Ethernet service network • Might require additional Ethernet service network switches or routers to accommodate new InfiniBand switches
Adding servers and a subnet to an existing InfiniBand network	• Connect the new InfiniBand switches to the Ethernet service network • Connect the new servers to Ethernet service network • Ensure operating system distribution capability for new servers without removable media • Might require additional HMCs to accommodate the new servers. If expanding beyond a single HMC or adding a CSM management server and cluster ready hardware server, you will need to deconfigure the current IP addressing method and reconfigure the network using DHCP on the CSM management server. • Might require additional Ethernet service network switches or routers to accommodate new InfiniBand switches

Management subsystem installation responsibilities:

Customer responsibilities are:

- The Ethernet service network
- HMCs
- CSM

- Customer installable CSM management server
- AIX NIM Spot installation to provide eServer diagnostics for servers without removable media.
- AIX NIM SPoT or Linux distribution server for distributing operating system updates to the servers without removable media (or to allow parallel updates to those with removable media).

IBM Service is responsible for installing the CSM management server.

Reference documentation for the management subsystem installation process:

The following documentation is needed for reference when you install the InfiniBand network management subsystem:

- HMC installation information
- CSM planning and installation guides
- Server installation guides for the CSM management server, AIX NIM server, and Linux distribution server
- NIM installation from AIX documentation from IBM System p and AIX Information Center Web site at http://publib.boulder.ibm.com/infocenter/pseries  .
- Linux distribution documentation

Management subsystem installation tasks:

During management subsystem install and configuration, you will be performing the following tasks:

1. Physically place the units on the floor.
2. Install and configure CSM and a cluster-ready hardware server.

   Note: This is only required in a cluster with multiple HMCs. If adding CSM MS and cluster-ready hardware server to an existing cluster, deconfigure current IP addressing and DHCP service on the Ethernet service network to accommodate the CSM MS as the DHCP server.
3. Install and configure the HMCs.
4. Build an AIX NIM SPoT for eServer diagnostics. This is only required if partitions do not have removable media (CD or DVD).
5. Build an AIX NIM SPoT or Linux distribution server for distributing operating system updates to the partitions without removable media (or to allow parallel updates to those with removable media).
6. Install the Ethernet service network.
7. Cable the management consoles and the Ethernet service network devices. Do not attach servers or switches to the Ethernet service network at this time.

Management subsystem installation process:

This process describes the steps that are required to install a new cluster management subsystem or to expand your existing cluster management subsystem.

1. M1 Physically place the management consoles (For example, HMCs, CSM management server, and NIM servers) and the Ethernet service network devices (switches and routers) on the floor.

   Note: HMCs might have maximum distance restrictions from the devices that they manage. To manage the new servers in your network, you might need to add one or more HMCs.
2. M2 You might choose to first install the Ethernet devices for the Ethernet service network. If you choose to do so, go to step 5 on page 44, and then return to step 3.
3. Choose from the following items, then go to the appropriate step for your cluster:

- If you are installing CSM with a Cluster Ready Hardware Server (this can be because it is required with multiple HMCs, or because the customer has opted to use CSM in the cluster), go to step 5.
- If you are installing a cluster with a single HMC without the cluster-ready hardware server, go to step 4.

4. To perform the installation of the management subsystem with a single HMC and without CSM and a cluster-ready hardware server, do the following:
   a. M2 Perform the HMC installation procedures:
      1) When you are installing and configuring the HMC, use the HMC documentation in the IBM Systems Hardware Information Center at http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp. See the Managing hardware platform consoles and interfaces topic collection.
      2) Ensure that the HMC is at the correct software and firmware levels. See the *README for IBM Clusters with the InfiniBand Switch* link at http://www14.software.ibm.com/webapp/set2/sas/f/networkmanager/home.html for information regarding the most current release levels of the HMC service update. Follow the links in the readme file to the appropriate download sites and instructions. Further information about HMC service level updates is available at http://www14.software.ibm.com/webapp/set2/sas/f/hmc/home.html. Be sure to follow the links for the appropriate level of HMC code.

         Note: IBM Network Manager (IBM NM) might have an additional service level (update.zip) to be applied after the HMC service level update is applied. After downloading the HMC service level update, return to the HMC support page for your HMC level, and click the **HPSNM/IBMNM** fixes tab for the latest IBM NM service level update. If a newer service level update is available, follow the instructions on the support page to download the latest service level update.
      3) If you downloaded an additional IBM NM service level update in the previous step, apply it now using the installation information found on the HMC support page.
   b. M2 Open a firewall port to allow SNMP communication between the IBM NM and the switches. For each HMC that might run IBM NM to manage the InfiniBand network, perform the procedure in "Adjusting Firewall Parameters for SNMP Traps" on page 146.
   c. M2 If you plan to have servers or partitions with no removable media (CD or DVD), build an AIX NIM SPoT on your chosen server to enable eServer diagnostics. Refer to NIM information in the AIX documentation.

      Note: Because the eServer diagnostics are available only in AIX, you will need an AIX SPoT, even if you are running another operating system in your partitions.
   d. M2 If you have servers with no removable media (CD or DVD), and you are going to use Linux in your partitions, install a Linux distribution server.
   e. Go to 6 on page 46. If you have already performed step 6 on page 46, proceed to step 7 on page 46.

5. To perform the installation of the management subsystem with multiple HMCs or CSM, and a cluster-ready hardware server, perform the following:

   Note: Perform this procedure if you are:
   - installing a new cluster with CSM and a cluster-ready hardware server.
   - adding an HMC to a cluster that already has CSM and a cluster-ready hardware server.
   - adding an HMC to a cluster with only a single HMC.
   - adding an InfiniBand network to an existing cluster with multiple HMCs that is not currently using CSM and cluster-ready hardware server.

a. Disable the DHCP server on the HMC and assign the HMC a static IP address so that there is only one DHCP server on the Ethernet service network, and so that device discovery will occur from the cluster-ready hardware server in CSM.

   Note: Disabling DHCP on the HMC will temporarily disconnect the HMC from its managed devices. If the current cluster already has CSM MS and a cluster-ready hardware server or does not require an additional HMC, go to step 5b.

   1) To perform a concurrent change of the DHCP server from the HMC to the CSM MS on a cluster server with 24-inch frames and BPCs (59*x* and 575 servers), the power subsystem microcode must be at level BP240_190, or later. When the DHCP server is restarted with level BP240_190 microcode, the BPCs will reacquire new ip addresses in the same manner as the service processors in the cluster servers.

      Note: If your 24-inch frame servers do not have the power subsystem microcode at level BP240_190, or later, your only option is to power off all BPCs (EPO-off) before moving the DHCP server from the HMC to the CSM MS.

   2) To enable the cluster-ready hardware server with CSM to connect correctly to the service processors and BPCs, set their authentication passwords to `abc123`. This is done using the HMC GUI using the following windows:
      - Frame Management Window for BPCs
      - Server and Partition Window for service processors

   3) Change existing HMCs from DHCP server to `static ip address` such that the address is within the cluster's Ethernet service network subnet (provided by the customer) but outside of the DHCP address range.
   4) Shutdown the HMC.
   5) Remove all BPCs and HMCs from the single system network, and continue as if this were a new install.

b. M2 Perform the initial installation and configuration of the HMCs using the HMC documentation, further details are available in the IBM Systems Hardware Information Center: http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp, See the Managing hardware platform consoles and interfaces topic collection.

   1) If there are multiple HMCs in the cluster or if you add an HMC to an existing cluster, you must use the CSM management server as the DHCP server. In this configuration, you must disable any HMC as a DHCP server and assign it a static IP address so that there is only one DHCP server on the Ethernet service network, and so that device discovery occurs from the cluster-ready hardware server in CSM.

      Note: Disabling DHCP on the HMC will temporarily disconnect the HMC from its managed devices. If the current cluster already has CSM MS and cluster-ready hardware server or does not require an additional HMC, go to step 5b.

   2) Ensure that your HMCs are at the correct software and firmware levels. See the *README for IBM Clusters with the InfiniBand Switch* link at http://www14.software.ibm.com/webapp/set2/sas/f/networkmanager/home.html for information regarding the most current released level of the HMC. Follow the links in the readme file to the appropriate download sites and instructions. Further information about HMC service level is available at http://www14.software.ibm.com/webapp/set2/sas/f/hmc/home.html. Follow the links for the appropriate level of HMC code.

      Note: IBM Network Manager (IBM NM) might have an additional service level (update.zip) to be applied after the HMC service level update is applied. After downloading the HMC service level update, return to the HMC support page for your HMC level, and click the **HPSNM/IBMNM fixes** tab for the latest IBM

NM service level update. If a service level update is available, follow the instructions on the support page to download the latest service level update.

3) If you downloaded an additional IBM NM service level update in the previous step, apply it now using the installation information found on the HMC support page.

c. M2 Open a firewall port to allow SNMP communication between the IBM NM and the switches. For each HMC that might run IBM NM to manage the InfiniBand network, perform the procedure in "Adjusting Firewall Parameters for SNMP Traps" on page 146.

Note: IBM Network Manager will only be active at any given time on one HMC in the cluster. However, if you have multiple HMCs, it is a good idea to enable the SNMP traps on at least two, if not all HMCs so that you can start IBM Network Manager on a backup HMC, if the primary HMC fails.

d. M2 Perform the procedures in the CSM installation guide. When performing those procedures, you must ensure that you do the following:

1) Install the CSM management server.
2) Update the operating system on the CSM management server.
3) Install the CSM code on the CSM management server.
4) Enable the CSM management server as the DHCP server for the Ethernet service network.
5) Define the subnet ranges for the Ethernet service networks.
6) Configure the DHCP ranges for the servers and BPCs.
7) Add the static IP addresses for the InfiniBand switches and HMCs to the Cluster Ready Hardware Server peer domain.
8) Start the DHCP server on the CSM Management Server.

e. M2 If you plan to have servers with no removable media (CD or DVD), build an AIX NIM SPoT on your chosen server to enable eServer diagnostics. Refer to NIM information in the AIX documentation.

Note: Because the eServer diagnostics are available only in AIX, you will need an AIX SPoT, even if you are running another operating system in your partitions.

f. M2 If you have servers with no removable media (CD or DVD), and you are going to use Linux in your partitions, install a Linux distribution server.

g. Go to step 6. If you have already performed step 6, then you may proceed to step 7.

6. M2 Install and configure the Ethernet service network using the documentation for the Ethernet devices and any configuration details provided by the HMC installation information. If you bypassed the previous steps to perform this step, return to step 3 on page 43.

7. M3 Cable the management consoles (HMCs and CSM Management Server) to the Ethernet service network.

Note: This step is only necessary for a new cluster installation, or if an expansion requires additional management consoles.

8. M4 Final configuration of management consoles: This task is performed in "Installing and configuring the cluster servers" on page 47 during the steps associated with S3 and M4. The following information is a preview of the tasks done in that procedure.

Note:

a. The BPCs and servers must be at power standby state before proceeding. See "Installing and configuring the cluster servers" on page 47 procedure up to and including major task S2.

b. DHCP (either the HMC or the CSM management server) must be up and running.

The following tasks are performed when you do the "Installing and configuring the cluster servers" procedure.

Note: Perform these tasks when servers and HCAs are added to your cluster.

- Verify that the BPCs and server service processors are connected to the DHCP server on the CSM management server.
- Setup the peer domains and HMC links in Cluster Ready Hardware Server on the CSM MS as instructed in the *CSM Planning and Install Guide*.
- Perform server and frame authentication with Cluster Ready Hardware Server on the CSM MS as instructed in the *CSM Planning and Install Guide*.

9. This procedure ends here. If you have responsibility for "Installing and configuring the cluster servers," proceed to that procedure. Otherwise, you may return to the overview in the "Installing a cluster that has an InfiniBand network" on page 39 section to find your next set of installation tasks.

Installing and configuring the cluster servers

Server installation and configuration encompasses major tasks **S1** - **S7**, and the server part of **M3** and **M4**, all of which are illustrated in Figure 3 on page 15. You will be installing and configuring the cluster's servers.

Note: If possible, do not begin this procedure until the "Install and configure the management subsystem" on page 41 procedure is completed. This helps alleviate the situation where installation personnel might be waiting on site for completion of key parts of this procedure. Review the Figure 3 on page 15 to identify the merge points where a step in a major task or procedure that is being performed by one person is dependent on the completion of steps in another major task or procedure that is being performed by another person.

Server installation and configuration responsibilities:

Customer responsibilities are:

- Installation of customer installable servers. This is dependent on which server models have been ordered.
- Updating system and power firmware.
- LPAR and HCA configuration.
- Loading and updating the operating systems.

IBM Service responsibilities are:

- Installation of servers that are IBM service representative installable servers. This is dependent on which servers models have been ordered.

Reference documentation for server installation and configuration procedure:

The following documentation is needed for reference when you perform these procedures:

- Server installation documentation
- HCA installation topic collections from the IBM Systems Hardware Information Center

Server installation and configuration tasks:

During server installation and configuration, you will perform the following tasks:

1. Place frames or racks and servers on the floor.
2. Start and verify servers and BPCs on the service network one frame or rack at a time.
3. Connect the cluster servers to their managing HMCs.
4. Authenticate frames and servers.
5. Update system and power firmware
6. Configure LPARs (including HCA configuration in the partition).

7. Verify the operation of the servers.
8. Install the operating systems.

Server installation and configuration procedure:

1. **S1** Position the frames or racks according to the floor plan.
2. Choose from the following items, then go to the appropriate step for your cluster:
   - If you have a single HMC in the cluster and you are not using CSM and a cluster-ready hardware server in your cluster, go to 3
   - If you are using CSM and cluster-ready hardware server in your cluster, go to 4.
3. If you have a single HMC and you are not using CSM and a cluster-ready hardware server in your cluster, do the following:
   a. **S1** Position the servers in frames or racks and install the host channel adapters (HCAs), do not connect or apply power to the servers at this time.

      Note: Do not proceed in the server installation instructions (WCII or Information Center) past the point where you physically install the hardware.

      Follow the installation procedures for servers found in:
      - Worldwide Customized Installation Instructions (WCIIs) for each server model that is installed by IBM service representatives.
      - For all other server models, customer procedures for initial server setup are available in the IBM Systems Hardware Information Center, http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp. Click **IBM Systems Hardware Information Center** → **Initial server setup**. Procedures for installing the GX host channel adapters are also available in the IBM Systems Hardware Information Center, click **IBM Systems Hardware Information Center** → **Installing hardware**.

   b. Verify that the HMC is running.
   c. After the Ethernet service network and management consoles have completed the initial installation and configuration, they are ready to discover and connect to frames and servers on the Ethernet service network. Proceed to step 5.
4. If you are using CSM and a cluster-ready hardware server in your cluster, do the following:
   a. **S1** Position servers in frames or racks and install the HCAs, do not connect or apply power to the servers at this time.

      Note: Do not proceed in the server install instructions (WCII or Information Center) past the point where you physically install the hardware.

      Follow the installation procedures for servers found in
      - Worldwide Customized Installation Instructions (WCIIs) for each server model that is installed by IBM service representatives.
      - For all other server models, customer procedures for initial server setup are available in the IBM Systems Hardware Information Center, http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp. Click **IBM Systems Hardware Information Center** → **Initial server setup**. Procedures for installing the GX host channel adapters are also available in the IBM Systems Hardware Information Center, click **IBM Systems Hardware Information Center** → **Installing hardware**.

   b. **S2** Verify that the DHCP server is running on the CSM management server.
   c. After the Ethernet service network and management consoles have completed the initial installation and configuration, they are ready to discover and connect to frames and servers on the Ethernet service network. Proceed to step 5.
5. For each server, at the HMC, set up the LPAR profiles.

Note: When setting up the LPAR profiles, you must configure the HCAs using the procedure found in "Installing or replacing a GX Host Channel Adapter" on page 64. Ensure that you do the step that configures the GUID index and capability for the HCA in the LPAR.

6. To bring the resources in each rack of servers onto the Ethernet service network and verify that addresses have been correctly served for each frame or rack of servers, perform the following procedure. By doing this one frame or rack at a time, you can verify that addresses have been served correctly, which is critical for cluster operation.
   a. M3 Connect the frame or server to the Ethernet service network. Use the documentation provided for the installation of these units.

      IBM Service personnel can access the WCIIs for each server model that is not a customer setup model. Customer server setup information is available in the IBM Systems Hardware Information Center, http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp. Click **IBM Systems Hardware Information Center → Initial server setup, and Installing hardware**

      Note: Do not proceed in the server install instructions (WCII or Information Center) past the point where you attach the Ethernet cables from the frames and servers to the Ethernet service network.

   b. Attach power cables to the frames and servers. Use the documentation provided for the installation of these units. For units that are installed by IBM Service, the service representative has access to WCIIs for each server model. For customer installable units, setup information is available in the IBM Systems Hardware Information Center, http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp. Click **IBM Systems Hardware Information Center -> Initial server setup, and Installing hardware**.

      Note: Do not power on the servers, continue with step 6c.

   c. S2 Apply power to the servers by connecting the power cords to a power source. If your servers are in frames with BPAs, press the EPO button for the frame. After the power is applied, the servers will power on and stop at the power standby state (not LPAR standby). For servers in frames or racks without BPAs, the server boots to the power standby state after connecting the power cable.

      Note: Do not press the power button on the control panels or apply power to the servers such that they boot to the LPAR standby state.

   d. S3 Use the following procedure to verify that the servers are now visible on the CSM management server.
      1) Check DHCP to verify that each server and BPC has been given an IP address by using the **dadmin –s** command.

         For a frame with a BPC, you should see an IP address assigned for each BPC and service processor connection. For a frame with no BPC, you will see IP addresses assigned for each service processor connection.

      2) Record the association between each server and its assigned IP address.
7. M4 After each server and BPC is visible on the CSM management server, using instructions for Cluster Ready Hardware Server in the CSM installation documentation, you must:
   a. Connect the frames and servers by assigning them to their respective managing HMC.
   b. Authenticate the frames and servers.
8. S3 In the server and frame management windows on each HMC, verify that you can see all the servers and frames to be managed by the HMC.
9. S4 Ensure that the servers and power subsystems (applies to IBM Systems with 24-inch racks) in your cluster are all at the correct software and firmware levels. See the *README*

for IBM Clusters with the InfiniBand switch file at http://www14.software.ibm.com/webapp/set2/sas/f/networkmanager/home.html for information regarding the most current release levels of:

- POWER™5 system firmware
- Power subsystem firmware (applies to IBM Systems with 24-inch racks)

Follow the links in the README file to the appropriate download sites and instructions.

Note: IBM Network Manager (IBM NM) might have an additional service pack (update.zip) to be applied after the HMC service pack update is applied. After downloading the HMC service pack, return to the HMC support page for your HMC level, and click the `HPSNM/IBMNM fixes` tab for the latest IBM NM service pack. If a service pack is available, follow the instructions on the support page to download the service pack.

10. **S5** Verify system operation from the HMCs by performing the following procedure at each HMC for the cluster:
    a. Bring the servers to LPAR standby and verify the system's viability by waiting several minutes and checking Service Focal Point.

       If you cannot bring a server to LPAR Standby, or there is a serviceable event reported in Service Focal Point, perform the prescribed service procedure as found in the IBM Systems Hardware Information Center, http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp.

    b. To verify each server, use the following procedure to run the eServer diagnostics:
       1) Depending on the server and who is doing the installation, you may wish to run these diagnostics from the CD-ROM, AIX NIM SPoT, or concurrently from an installed AIX operating system. If you wish to run this from the AIX operating system on the server, you will have to wait until the LPARs are configured and the operating system is installed.
       2) To resolve any problem with a server, check the diagnostic results and Service Focal Point and follow the maintenance procedures.

       Note: Typically, the IBM service representative's responsibility ends here for IBM service installed frames and servers. From this point forward, after the IBM service representative leaves the site, if any problem is found in a server, or with an InfiniBand link, a service call must be placed.

       The IBM service representative should recognize that, at this point, the HCA link interface and InfiniBand cables have not been verified, and will not be verified until the end of the procedure for InfiniBand network verification, which may be performed by either the customer or a non-IBM vendor. When the IBM service representative leaves the site, it is possible that the procedure for InfiniBand network verification may identify a faulty link, in which case the IBM service representative may receive a service call to isolate to and repair the cable or HCA on the faulty link.
11. **S6** Customize LPARs and HCA configuration
    a. Define LPARs using the procedures found in IBM Systems Hardware Information Center, http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp.
    b. Configure the HCAs using the procedure found in "Installing or replacing a GX Host Channel Adapter" on page 64. Ensure that you do the steps that configure the GUID index and capability for the HCA in the LPAR.
12. **S7** Install and update the operating systems

    If servers do not have removable media, you must use an AIX NIM server, or Linux distribution server to load and update the operating systems.

13. If you have responsibility for "InfiniBand switch installation and configuration for vendor switches" on page 55, proceed to that procedure. Otherwise, return to the overview of the installation section to find your next set of installation tasks. **This procedure ends here.**

Server installation when expanding a cluster

If you are adding or expanding InfiniBand network capabilities to an existing cluster by adding servers to the cluster, then you might need to approach the Server installation and configuration a little differently than with a new cluster process. The process for server installation and configuration is based on a new cluster installation, but it will indicate where there are variances for expansion scenarios.

The following table outlines how the new cluster installation is affected/altered by expansion scenarios:

Table 14. Expanding a cluster by adding a seerver

Scenario	Effects
Adding InfiniBand hardware to an existing cluster (switches and HCAs)	Configure the LPARs to use the HCAs Configure HCAs for switch partitioning.
Adding new servers to an existing InfiniBand network	Perform this task as if it were a new cluster installation.
Adding HCAs to an existing InfiniBand network	Perform this task as if it were a new cluster installation.
Adding a subnet to an existing InfiniBand network	Configure the LPARs to use the new HCA ports. Configure the newly cabled HCA ports for switch partitioning.
Adding servers and a subnet to an existing InfiniBand network	Perform this task as if it were a new cluster installation.

Server installation and configuration responsibilities:

Customer responsibilities are:

- Installation of customer installable servers.
- Updating system and power firmware.
- LPAR and HCA configuration.
- Loading and updating the operating systems.

IBM Service responsibilities are:

- Installation of servers that are IBM service representative installable servers.

Reference documentation for server installation and configuration procedure:

The following documentation is needed as reference for these procedures:

- Server installation documentation
- HCA installation topic collections from the IBM Systems Hardware Information Center

Server installation and configuration tasks:

During server installation and configuration, you will perform the following tasks:

1. Place frames or racks and servers on the floor.
2. Start and verify servers and BPCs on the service network one frame or rack at a time.
3. Connect the cluster servers to their managing HMCs.
4. Authenticate frames and servers.
5. Update system and power firmware
6. Configure LPARs (including HCA configuration in the partition).
7. Verify the operation of the servers.
8. Install the operating systems.

Server installation and configuration procedure:

1. Select one of the following options:
   - If you are adding servers to an existing cluster, go to 2
   - If you are adding cables to existing HCAs, proceed to step 12 on page 54.
   - If you are adding HCAs to existing servers, go to "Installing or replacing a GX Host Channel Adapter" on page 64 and follow the installation instructions for the HCAs (WCII or Information Center instructions), then proceed to step 12 on page 54.
2. **S1** Position the frames or racks according to the floor plan.
3. Choose from the following items, then go to the appropriate step for your cluster:
   - If you have a single HMC in the cluster and you are not using CSM and a cluster-ready hardware server in your cluster, go to 4
   - If you are using CSM and cluster-ready hardware server in your cluster, go to 5.
4. If you have a single HMC and you are not using CSM and a cluster-ready hardware server in your cluster, do the following:
   a. **S1** Position the servers in frames or racks and install the host channel adapters (HCAs), do not connect or apply power to the servers at this time.

      Note: Do not proceed in the server installation instructions (WCII or Information Center) past the point where you physically install the hardware.

      Follow the installation procedures for servers found in:
      - Worldwide Customized Installation Instructions (WCIIs) for each server model that is installed by IBM service representatives.
      - For all other server models, customer procedures for initial server setup are available in the IBM Systems Hardware Information Center, http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp. Click **IBM Systems Hardware Information Center → Initial server setup**. Procedures for installing the GX host channel adapters are also available in the IBM Systems Hardware Information Center, click **IBM Systems Hardware Information Center → Installing hardware**.

   b. Verify that the HMC is running.
   c. After the Ethernet service network and management consoles have completed the initial installation and configuration, they are ready to discover and connect to frames and servers on the Ethernet service network. Proceed to step 6 on page 53.
5. If you are using CSM and a cluster-ready hardware server in your cluster, do the following:
   a. **S1** Position servers in frames or racks and install the HCAs, do not connect or apply power to the servers at this time.

      Note: Do not proceed in the server install instructions (WCII or Information Center) past the point where you physically install the hardware.

      Follow the installation procedures for servers found in
      - Worldwide Customized Installation Instructions (WCIIs) for each server model that is installed by IBM service representatives.
      - For all other server models, customer procedures for initial server setup are available in the IBM Systems Hardware Information Center, http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp. Click **IBM Systems Hardware Information Center → Initial server setup**. Procedures for installing the GX host channel adapters are also available in the IBM Systems Hardware Information Center, click **IBM Systems Hardware Information Center → Installing hardware**.

   b. **S2** Verify that the DHCP server is running on the CSM management server.
   c. After the Ethernet service network and management consoles have completed the initial installation and configuration, they are ready to discover and connect to frames and servers on the Ethernet service network. Proceed to step 6 on page 53.

6. For each server, at the HMC, set up the LPAR profiles.

   Note: When setting up the LPAR profiles, you must configure the HCAs using the procedure found in "Installing or replacing a GX Host Channel Adapter" on page 64. Ensure that you do the step that configures the GUID index and capability for the HCA in the LPAR.

7. To connect the resources in each rack of servers to the Ethernet service network and verify that addresses have been correctly served for each frame or rack of servers, perform the following procedure. By doing this one frame or rack at a time, you can verify that addresses have been served correctly, which is critical for cluster operation.

   a. M3 Connect the frame or server to the Ethernet service network. Use the documentation provided for the installation of these units.

      IBM Service personnel can access the WCIIs for each server model that is not a customer setup model. Customer server setup information is available in the IBM Systems Hardware Information Center, http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp. Click **IBM Systems Hardware Information Center → Initial server setup, and Installing hardware**

      Note: Do not proceed in the server install instructions (WCII or Information Center) past the point where you attach the Ethernet cables from the frames and servers to the Ethernet service network.

   b. Attach power cables to the frames and servers. Use the documentation provided for the installation of these units. For units that are installed by IBM Service, the service representative has access to WCIIs for each server model. For customer installable units, setup information is available in the IBM Systems Hardware Information Center, http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp. Click **IBM Systems Hardware Information Center -> Initial server setup, and Installing hardware**.

      Note: Do not power on the servers, continue with step 7c.

   c. S2 Apply power to the servers by connecting the power cords to a power source. If your servers are in frames with BPAs, press the EPO button for the frame. After the power is applied, the servers will power on and stop at the power standby state (not LPAR standby). For servers in frames or racks without BPAs, the server boots to the power standby state after connecting the power cable.

      Note: Do not press the power button on the control panels or apply power to the servers such that they boot to the LPAR standby state.

   d. S3 Use the following procedure to verify that the servers are now visible on the CSM management server.

      1) Check DHCP to verify that each server and BPC has been given an IP address by using the **dadmin –s** command.

         For a frame with a BPC, you should see an IP address assigned for each BPC and service processor connection. For a frame or rack with no BPC, you will see IP addresses assigned for each service processor connection.

      2) Record the association between each server and its assigned IP address.

8. M4 After each server and BPC is visible on the CSM management server, using instructions for Cluster Ready Hardware Server in the CSM installation documentation, you must:

   a. Connect the frames and servers by assigning them to their respective managing HMC.

   b. Authenticate the frames and servers.

9. S3 In the server and frame management windows on each HMC, verify that you can see all the servers and frames to be managed by the HMC.

10. **S4** Ensure that the servers and power subsystems (applies to IBM Systems with 24-inch racks) in your cluster are all at the correct software and firmware levels. See the *README for IBM Clusters with the InfiniBand switch* file at http://www14.software.ibm.com/webapp/set2/sas/f/networkmanager/home.html for information regarding the most current release levels of:
    - POWER5 system firmware
    - Power subsystem firmware (applies to IBM Systems with 24-inch racks)

    Follow the links in the README file to the appropriate download sites and instructions.

    Note: IBM Network Manager (IBM NM) might have an additional service pack (update.zip) to be applied after the HMC service pack update is applied. After downloading the HMC service pack, return to the HMC support page for your HMC level, and click the `HPSNM/IBMNM fixes` tab for the latest IBM NM service pack. If a service pack is available, follow the instructions on the support page to download the service pack.
11. **S5** Verify system operation from the HMCs by performing the following procedure at each HMC for the cluster:
    a. Bring the servers to LPAR standby and verify the system's viability by waiting several minutes and checking Service Focal Point.

       If you cannot bring a server to LPAR Standby, or there is a serviceable event reported in Service Focal Point, perform the prescribed service procedure as found in the IBM Systems Hardware Information Center, http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp.
    b. To verify each server, use the following procedure to run the eServer diagnostics:
        1) Depending on the server and who is doing the installation, you may want to run these diagnostics from the CD-ROM, AIX NIM SPoT, or concurrently from an installed AIX operating system. If you want to run this from the AIX operating system on the server, you will have to wait until the LPARs are configured and the operating system is installed.
        2) To resolve any problem with a server, check the diagnostic results and Service Focal Point and follow the maintenance procedures.

        Note: Typically, the IBM service representative's responsibility ends here for IBM service installed frames and servers. From this point forward, after the IBM service representative leaves the site, if any problem is found in a server, or with an InfiniBand link, a service call must be placed.

        The IBM service representative should recognize that, at this point, the HCA link interface and InfiniBand cables have not been verified, and will not be verified until the end of the procedure for InfiniBand network verification, which may be performed by either the customer or a non-IBM vendor. When the IBM service representative leaves the site, it is possible that the procedure for InfiniBand network verification may identify a faulty link, in which case the IBM service representative might receive a service call to isolate to and repair the cable or HCA on the faulty link.
12. **S6** Customize LPARs and HCA configuration
    a. Define LPARs using the procedures found in IBM Systems Hardware Information Center, http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp.
    b. Configure the HCAs using the procedure found in "Installing or replacing a GX Host Channel Adapter" on page 64. Ensure that you do the steps that configure the GUID index and capability for the HCA in the LPAR.
13. **S7** Install and update the operating systems

    If servers do not have removable media, you must use an AIX NIM server, or Linux distribution server to load and update the operating systems.

14. If you have responsibility for "InfiniBand switch installation and configuration for vendor switches," proceed to that procedure. Otherwise, return to the overview of the installation section to find your next set of installation tasks. **This procedure ends here.**

InfiniBand switch installation and configuration for vendor switches

The InfiniBand switch installation and configuration encompasses major tasks **W1** - **W6**, which are illustrated in the Figure 3 on page 15 figure.

Note: If possible, this procedure should not begin before the management subsystem installation and configuration procedure is completed. This will alleviate the situation where various installation personnel may be waiting on site for key parts of this procedure to be completed. Depending on the arrival of units on site, this is not always practical. Therefore, it is important to review the "Order of Installation" on page 13 and Figure 3 on page 15 in "Planning for InfiniBand networks" on page 5 to identify the merge points where a step in a major task or procedure being performed by one person is dependent on the completion of steps in another major task or procedure being performed by another person.

InfiniBand switch installation and configuration information for expansion

If you are adding or expanding InfiniBand network capabilities to an existing cluster, then you may need to approach the InfiniBand switch installation and configuration a little differently than with a new cluster flow. The flow for InfiniBand switch installation and configuration is based on a new cluster installation, but it will indicate where there are variances for expansion scenarios.

The following table outlines how the new cluster installation is affected by expansion scenarios:

Table 15.

Scenario	Effects
Adding InfiniBand hardware to an existing cluster (switches and HCAs)	Perform this task as if it were a new cluster installation.
Adding new servers to an existing InfiniBand network	If you need to add an HMC to a cluster that only had a single HMC and no CSM MS with cluster-ready hardware server, you will need to reconfigure the existing switches to use static-IP addressing on the Ethernet service network.
Adding HCAs to an existing InfiniBand network	You should not have to perform anything outlined in this major task.
Adding a subnet to an existing InfiniBand network	Perform this task on new switches as if it were a new cluster installation.
Adding servers and a subnet to an existing InfiniBand network	Perform this task on new switches as if it were a new cluster installation. If you need to add an HMC to a cluster that only had a single HMC and no CSM MS with cluster-ready hardware server, you will need to reconfigure the existing switches to use static-IP addressing on the service Ethernet network.

InfiniBand switch installation and configuration responsibilities

For switches that are not 7048-120 or 7048-270, the customer or the switch vendor is responsible for installing the InfiniBand switches. If the switch is a 7048-120 or 7048-270, IBM is responsible for the installation.

The customer is responsible for verifying that the switch is detected by IBM Network Manager.

Reference documentation for InfiniBand switch installation and configuration procedure

While you perform this procedure, you will need the InfiniBand switch vendor's installation documentation for reference.

InfiniBand switch installation and configuration tasks:

During switch install and configuration, you will perform the following tasks:

Note: If you are expanding an existing cluster with an additional switch, the installation process is similar to installing a new cluster. However, if the new configuration changes the DHCP support from an HMC to a separate cluster-ready hardware server, additional steps must be taken during the installation of the switch.

1. Place frames or racks and switches on the floor.
2. Power on InfiniBand switches and configure their IP addresses on the service Ethernet network. (If DHCP support moves to a cluster-ready hardware server this task requires some manual fixed IP address updates.)
3. Connect the new InfiniBand switches to the service Ethernet network
4. Verify InfiniBand switches are discovered on the Ethernet service network
5. Update InfiniBand switch software
6. Customize InfiniBand Network Configuration (including GID-prefixes and LMC values)

InfiniBand switch installation and configuration procedure

It is possible to perform some of the tasks in this procedure in a method other than that which is described. If you have other methods for configuring switches, you also must review a few key points in the installation process with regard to order and coordination of tasks and configuration settings that are required in a cluster environment. Review the following list of key points before beginning the switch installation process:

- Power on the InfiniBand switches and configure their IP addresses before attaching them to the Ethernet service network.
- If an InfiniBand switch has multiple Ethernet connections for the Ethernet service network, and the cluster has multiple service Ethernet subnetworks, the switch's Ethernet ports must connect to the same subnetwork.
- In a cluster with a single HMC that does not use CSM and Cluster Ready Hardware Server, you must configure the switch Ethernet IP addressing for DHCP so that the HMC, acting as the DHCP server, may assign the address.
- In a cluster with multiple HMCs or one that uses CSM and Cluster Ready Hardware Server, you must configure the switch Ethernet to a static address which is provided by the customer.
- Set the GID-prefix value according to the installation plan.
- If this is an HPC environment, set the LMC value to 2.
- If there is a single switch in a subnet, disable database synchronization on that switch.
- If there are multiple switches on a subnet, enable database synchronization on all switches in the subnet.
- Update the switch firmware code as required. See the *README for IBM Clusters with the InfiniBand switch* link at http://www14.software.ibm.com/webapp/set2/sas/f/networkmanager/home.html for information regarding switch code levels.
- Instruct the customer to verify that the switch is detected by IBM Network Manager using the verify detection step in the following procedure:

If you are expanding an existing cluster, also consider the following items:

- For new InfiniBand switches, perform all the steps in the following procedure on the new InfiniBand switches.
- If you need to add an HMC or CSM and cluster-ready hardware server to an existing network, go to step 2.

Do the following procedure to install and configure your InfiniBand switches:

1. Physically place frames and switches on the floor:
   a. Review the vendor documentation for each switch model that you are installing.
   b. Physically install InfiniBand switches into 19-inch frames (or racks) and attach power cables to the switches according to the instructions for the InfiniBand switch model. This will automatically power on the switches. There is no power switch for the Topspin or Cisco switch models.

      Note: Do not connect the Ethernet connections for the service network at this time.
   c. For any Cisco 7008 Server Switch or SFS7008P switch that has fewer than 8 line interface modules (LIMs) be sure to record the slots in which the LIMs are populated and the slots in which the fabric controller cards are populated. This is done to ensure that the LIMs are populated in a manner that allows for good cable management and expansion capability. Also, this ensures that the fabric controllers are populated in a manner that will support the LIMs that are populated.
      - The LIMs are in slots 1 through 8 in the back of the switch, and they have the cable connectors.
      - The Fabric Controller cards are in slots 9 through 14 in the front of the switch. See "Location codes used by the IBM Network Manager" on page 218 for the 7048-270 or SFS7008P.
      - Ensure that any empty LIM slots are above the populated slots.
      - Ensure that the LIMs are populated in pairs and have corresponding node fabric controller cards as described in "Plugging Rules for 7048-270 or SFS7008P switch line interface modules and fabric controllers" on page 152.
2. Set up the Ethernet interface for the service network:

   Note: If you are adding an HMC to a cluster that had only a single HMC or if you are adding CSM with cluster-ready hardware server to a cluster that did not have it previously, you must perform step 2b. Existing switches must be rebooted for this, and thus there will be an interruption to data passing on the InfiniBand network.
   a. If this is a single HMC configuration, enable the DHCP client on the switch and save the configuration. Use the procedure for setting the switch up as a DHCP client in "Setting IP addressing in switches" on page 159.
   b. If there are multiple HMCs in this configuration and the CSM Management Server is setup, set the switch to a fixed IP-address provided by the customer and save the configuration. Use the procedure for setting the switch up with a static IP address in "Setting IP addressing in switches" on page 159.

      Note: Because you are already in the Command Line Interface (CLI) during this step, you should review the next step and see if you need to alter the database synchronization setting before exiting from the CLI.
3. The database synchronization value needs to be correct for your configuration. Use the procedure "Setting database synchronization" on page 161 to ensure the value is correctly set for your switch configuration.
4. Attach the switch to the Ethernet service network.

   Note: If the switch has multiple Ethernet connections, they must all attach to the same subnet.
5. Verify that the installed switch is detected by IBM Network Manager by using the Management Properties View:

Note: If you are expanding a cluster with a new switch, perform this procedure for each new switch added, and for existing switches if they have had their service IP addressing method changed between DHCP and static.

a. Go to the HMC that is running IBM Network Manager. If no HMC is currently running IBM Network Manager, go to the HMC that is designated for running IBM Network Manager and enable IBM Network Manager:
   - On the HMC select **Switch Management**.
   - Select **IBM Network Manager**.
   - Select **Enable IBM NM Software**.

b. From the IBM Network Manager Overview window, click **View Management Properties**.

c. Click the **Switch** tab.

   Note: If a switch is not listed, investigate the Ethernet service network connections through to the switch. If these connections appear operational, consider that the IP addressing on the switch might not be configured correctly. If there is more than one HMC in the configuration, the switch must be configured for static-IP on its Ethernet interface. If there is only one HMC in the configuration, DHCP must be used.

d. Verify that all switches are listed in the Management Properties window.

e. Verify that all switches have a valid location code. Valid Formats are:
   - U7048.120.[serial number]-P1 for a 7048-120
   - U7048.270.[serial number]-P1 for a 7048-270
   - USFS7000Pxxx.[serial number]-P1 for an SFS7000P (where xxx = a variable length string)
   - USFS7008Pxxx.[serial number]-P1 for an SFS7008P (where xxx = a variable length string) If a switch does not have a valid location code format, it is not supported in a cluster. It is also possible that the VPD module is corrupted. Check the labeling on the switch to verify that it is the correct model. If the label indicates an incorrect model, call your next level of support, because the switch is not supported in the cluster. Otherwise:
     – If it is a SFS7000P switch, have the switch replaced.
     – If it is a SFS7008P switch, replace the chassis-ID module and power on the switch by unplugging and plugging the cables.

6. Verify that the switch code matches the latest supported level using the procedure in “Checking switch software levels” on page 144.

   Note: If you are adding a switch to an existing cluster, refer to InfiniBand switch vendor documentation to determine if you should make sure that all switches (new and existing) are at the same switch code level, or if there is compatibility between levels.

   a. If the software level is incorrect, have the customer obtain the latest level from the Cisco or Topspin support site.

   b. Use the procedure found in “Updating switch software” on page 144.

7. Finalize the configuration for each InfiniBand switch:

   Note: If you are adding a switch to an existing cluster, this is required for only new switches, unless it has been decided that different GID-prefixes or LMCs are required for the existing switches.

   a. Start IBM Network Manager on an HMC.

   b. Set the GID-prefix for each switch using the GID-prefixes provided by the customer’s planning documentation and the procedure in “Setting GID prefixes” on page 147.

   c. If the switches are used for HPC applications, set the LMC for each switch to 2 using the procedure in “Setting the location identifier mask control” on page 148.

   d. If the switches are not for HPC applications, the default LMC of 0 is acceptable. To check the current setting use the procedure in “Checking the logical identifier mask control” on page 147.

e. Set the switch name.

8. This procedure ends here. If you are also responsible for cabling the InfiniBand network, proceed to "Attach cables to the InfiniBand network." Otherwise, you may return to the overview of the install section to find your next set of install tasks.

Attach cables to the InfiniBand network

Cabling the InfiniBand network encompasses major tasks **C1** - **C4** , which are illustrated in the Figure 3 on page 15 figure.

Note: Do not start this procedure until InfiniBand switches have been physically installed. Wait until the servers have been configured. This will alleviate the situation where various install personnel may be waiting on site for key parts of this procedure to be completed. Depending on the arrival of units on site, this is not always practical. Therefore, it is important to review the "Order of Installation" on page 13 and Figure 3 on page 15 figure from Planning for InfiniBand networks to identify the merge points where a step in a major task or procedure being performed by one person is dependent on the completion of steps in another major task or procedure being performed by another person.

Cabling the InfiniBand network information for expansion

If you are adding or expanding InfiniBand network capabilities to an existing cluster, then you may need to approach Cabling the InfiniBand a little differently than with a new cluster flow. The flow for Cabling the InfiniBand network is based on a new cluster installation, but it will indicate where there are variances for expansion scenarios. The following table outlines how the new cluster installation is affected/altered by expansion scenarios:

Table 16.

Header	Header
Adding InfiniBand hardware to an existing cluster (switches and HCAs)	Perform this task as if it were a new cluster installation. All InfiniBand hardware is new to the cluster.
Adding new servers to an existing InfiniBand network	Perform this task as if it were a new cluster installation for all new servers and HCAs added to the existing cluster.
Adding HCAs to an existing InfiniBand network	Perform this task as if it were a new cluster installation for all new HCAs added to the existing cluster.
Adding a subnet to an existing InfiniBand network	Perform this task as if it were a new cluster installation for all new switches added to the existing cluster.
Adding servers and a subnet to	Perform this task as if it were a new cluster installation for all new servers and HCAs and switches added to the existing cluster.

InfiniBand network cabling responsibilities:

The responsibilities for InfiniBand network cabling install may be the customer's responsibility or non-IBM vendor's responsibility.

The IBM service representative is responsible for replacing faulty or damaged cables. If the switch being installed is a 7048-120 or 7048-270, the IBM service representative is responsible for this cabling procedure.

Reference documentation for InfiniBand network cabling procedure:

This is the documentation that you will require as reference for this procedure: InfiniBand switch vendor install documentation.

InfiniBand network cabling tasks:

During server install and configuration, you will perform the following tasks:

1. Review the cable plan for the InfiniBand network.

2. Prepare the cable labels.
3. Install any octopus cables and configure 4x switch ports into groups to use them.
4. Install HCA to switch cables.
5. Install switch to switch cables.

InfiniBand network cabling procedure:

It is possible to perform some of the tasks in this procedure in a method other than that which is described. If you have other methods for cabling the InfiniBand network, you still must review a few key points in the installation process with regard to order and coordination of tasks and configuration settings that are required in a cluster environment:

- IBM is responsible for faulty or damaged cable replacement.
- If octopus cables are used, install them in the following order:
  1. You must adhere to the cable plan so that the 4x cable connectors of the 12x octopus cable connect to the proper switch port group. If this was not planned ahead of time, see "Planning octopus cables in static 12x cabling" on page 19, and "Planning 4x connections for octopus cables" on page 21.
  2. Connect the switch end of the cables. The configuration step that follows can be done before this step. However, you cannot connect the HCAs until both the configuration step and the connection to the switch end of the cables is done.
  3. Configure the lowest numbered port in a switch port group to an admin-speed of 30 Gbps and disable the auto-negotiate setting.
  4. Connect the 12x connector of the octopus cable to the HCA port.

     Note: When you connect an octopus cable from a switch to an HCA port, remember to connect 12x connector to the HCA port after connecting 4x connectors to the switch. If the correct order is not followed, problems might be induced in the HCA logic that determines the link speed.

Do the following to complete your switch network cabling:

1. Obtain and review a copy of the cable plan for the InfiniBand network.
2. Route the InfiniBand cables according to the cable plan and attach them to only the switch ports. Refer to the Cisco or Topspin documentation for more information on how to plug cables. If you do not have octopus cables to connect, proceed to step 4 on page 62
3. If you have any 12x HCAs that will be connected to 4x switches with octopus cables, you must do additional configuration setup and then plug the cables in the correct order. For each octopus cable, do the following:
   a. To configure the static-12x groups, you must do so through the switch Command Line Interface (CLI). Connecting to and logging onto the switch CLI is described in the switch installation instructions, see "Accessing an InfiniBand switch command line interface" on page 158 and the switch hardware guides. When you have the CLI prompt, perform the following:
   b. Log into the command with the Login: `super`
   c. Enter: `enable`
   d. Enter: `show interface ib [card]/[slot]`

      Note: card = 1 for a 7048-120 or 7048-700, or the LIM number for the 7048-270 or 7048-708. The slot is the lowest numbered port in the group, which you should have recorded when determining the groups.

      If you entered 1/1 for the [card]/[slot], the following output will be seen.

```
================================================================================
InfiniBand Interface Information
================================================================================
port : 1/1
name : 1/1
type : ib4xTX
desc : 1/1 (65)
last-change : Wed Dec 7 17:10:03 2005
mtu : 2048 auto-negotiate-supported : yes
auto-negotiate : disabled
admin-status : up
oper-status : up
admin-speed : 10gbps
oper-speed : unknown
phy-state : polling
link-trap : enabled
```

e. Record the value for the auto-negotiate parameter.

f. Enter: `config`.

g. Enter: `interface ib [card]/[slot]`.

   The prompt will change to indicate that you are configuring the card and slot. The following example is for card 1, slot 1: **Topspin-270(config-if-ib-1/1)**.

h. Enter: `no auto-negotiate` .

   Note: The `no` ensures that the auto-negotiate mode will be disabled.

i. Enter: `speed 12x`.

j. If you have more groups to configure, repeat steps 3g through 3i until you have configured each group.

k. Enter: `exit`.

l. To preserve changes across reboots enter: `copy running-config startup-config`.

m. Enter: `exit`.

   You are now out of the config mode in the CLI.

n. Enter: `show interface ib [card]/[slot]`. The output shows that the admin-speed is 30 gbps, and auto-negotiate is disabled. This step assumes that you have connected the cables, so no other parameters are significant at this point. Be very sure that auto-negotiate is disabled. If it is enabled you are not in static-12x mode, regardless of the value of admin-speed, or oper-speed. The following is an example output for card/slot = 1/1.

```
================================================================================
InfiniBand Interface Information
================================================================================
port : 1/1
name : 1/1
type : ib4xTX
desc : 1/1 (65)
last-change : Wed Dec 7 17:10:03 2005
mtu : 2048 auto-negotiate-supported : yes
auto-negotiate : disabled
admin-status : up
oper-status : up
admin-speed : 30gbps
oper-speed : 30gbps
phy-state : linkup
link-trap : enabled
```

o. If you configured more than one group, you can repeat step 3n for each group of switch ports.

p. To leave the CLI, Enter: exit

Note: Repeat the previous procedure for each switch and group of switch ports to which octopus cables connect.

4. Connect the InfiniBand cables to the switch ports according to the planning documentation. Refer to the Cisco or Topspin hardware guide for more information on how to plug cables. Do not plug the cables into the HCA ports, yet. Note: If you have octopus cables, the ordering of the 4x cable ends should have been defined during cable planning. It is very important to get the ordering of the three 4x cable connectors correct. If this was not done, refer to Planning the connection of the 4x connectors of the octopus cables, which will help you understand how to plug the cables. The connector labeled "16" sometimes connects to the lowest numbered port in a group, and sometimes to the highest numbered port in a group.
5. Connect the InfiniBand cables to the HCA ports according to the planning documentation.
6. This procedure ends here. If you are responsible for verifying the InfiniBand network topology and operation, you may proceed to that procedure. Otherwise, you may return to the overview of the install section to find your next set of install tasks.

Verify the InfiniBand network topology and operation

Verifying the InfiniBand network topology and operation encompasses major tasks **V1** - **V3**, which are illustrated in the Figure 3 on page 15 figure.

Note: This procedure cannot be performed until all other procedures in cluster installation have been completed. These include the management subsystem installation and configuration, server installation and configuration, InfiniBand switch installation and configuration, and attaching cables to the InfiniBand network.

Verifying InfiniBand network topology and operation responsibilities:

The responsibilities for InfiniBand network cabling can be the customer's responsibility or a non-IBM vendor's responsibility.

IBM service is responsible for replacing faulty or damaged IBM-supplied cables. If the switch being installed is a 7048-120, 7048-270, the IBM service representative is responsible for the verification of the topology.

Reference documentation for verifying InfiniBand network topology and operation procedure:

No special documentation is referenced in this procedure.

Verifying InfiniBand network topology and operation tasks:

During server install and configuration, you will perform the following tasks:

1. Verifying that the network topology and cabling is complete according to plan
2. Verifying the cluster network operation by passing data and checking for errors

Verifying InfiniBand network topology and operation Procedure:

It is possible to perform some of the tasks in this procedure in a method other than that which is described. If you have other methods for cabling the InfiniBand network, you still must review a few key points in the installation process with regard to order and coordination of tasks and configuration settings that are required in a cluster environment:

- This procedure cannot be performed until all other procedures in the cluster installation have been completed. These include the management subsystem installation and configuration, server installation and configuration, InfiniBand switch installation and configuration, and cabling the InfiniBand network.
- IBM service is responsible for faulty or damaged cable replacement.
- The customer should check the availability of HCAs to the operating system before any application is run to verify network operation.

- If octopus cables are used, verify that the cables are connected correctly by using "Verifying static-12x mode connectivity" on page 67. This procedure includes expected LED states and expected IBM Network Manager status for each port in a group of ports working as a single 12x link.
- If you find a problem with a link that might be caused by a faulty host channel adapter or cable, contact your service representative for repair.
- This is the final procedure in installing an IBM System p cluster with an InfiniBand network.

The following procedure provides additional details that can help you perform the verification of your network.

1. Do the following to verify the network topology:
   a. Check all LEDs for the switch ports to verify that they are properly lit. See "Interpreting LEDs" on page 122 for normal 4x cable connections. See "Verifying static-12x mode connectivity" on page 67 for octopus cable connections to switches.
   b. Go to the HMC that is running IBM Network Manager. If no HMC is currently running IBM Network Manager, go to the HMC that is designated for running IBM Network Manager and enable IBM Network Manager:
      1) On the HMC GUI click, **Switch Management**.
      2) Click **IBM Network Manager**.
      3) Click **Enable IBM NM Software**.

      Note: Only one instance of IBM Network Manager can be run in a cluster. Unpredictable results will occur if more than one instance is running.

      If you decide to move the IBM Network Manager functions from one HMC host to another HMC host, you must first disable the IBM Network Manager on the first HMC host before enabling it on the second host.
   c. From the IBM Network Manager Overview window, click **View Switch Topology**.
   d. All of the switches that are powered on should be displayed in the Switch Topology view. If there is a missing switch, verify that the power cables are still plugged in and that there is connectivity to the service network through to the HMC running IBM Network Manager. Further information on checking LEDs on the switch can be found in Switch chassis and system-wide LEDs, and in the respective Cisco or Topspin hardware service guides.
   e. Click **File → Expand All**.
   f. Do the following to verify the switch network configuration that is displayed:
      1) If you have installed octopus cables to attach 12x HCAs to 4x switches, use the appropriate verification procedures for your switch found in "Verifying static-12x mode connectivity" on page 67.
      2) In IBM Network Manager's Switch Topology view, verify that the neighbors for the switch ports are as expected according to the cabling plan. To understand location code formats, go to "InfiniBand component location codes" on page 217. If there are any cables that are connected to incorrect ports, correct the cabling.
      3) If a 7048-120 or SFS7000P is missing ports, try re-powering the switch. Otherwise, replace the switch.
      4) If a 7048-270 or SFS7008P is missing ports, it is possible that there is a problem with a LIM or a fabric controller. Check the following:
         a) Check the LEDs for the LIM and its corresponding Fabric controller. For information on LEDs, see "Interpreting LEDs" on page 122. For information on LIM and fabric controller relationships, see "Plugging Rules for 7048-270 or SFS7008P switch line interface modules and fabric controllers" on page 152 LIMs and fabric controllers.

b) Verify that the LIMs and associated fabric controllers are properly seated.

5) Check Service Focal Point (SFP) on all HMCs and make any required repairs. Contact IBM service to repair any serviceable events reported to SFP.

2. Do the following to verify cluster network operation:
   a. Verify that the HCAs are available to the operating system in each LPAR:
      1) For partitions running AIX, check the HCA status by running the lsdev -C | grep ib . An example of good results for verifying a GX HCA is: **iba0 Available** Host Channel Adapter.
      2) For partitions running Linux, refer to the documentation provided with SUSE Linux Enterprise Server 9 (SP2) with IBM IB GX HCA driver and OpenIB Gen3 Stack. Refer to the instructions in the eHCAD install file contained within the download from SourceForge web site, http://sourceforge.net/projects/ibmehcad.
   b. At this time, you should run an all-to-all ping test across the InfiniBand network.
   c. After running for a while, check Service Focal Point on all HMCs. If there is a serviceable event reported, contact IBM Service.
3. The InfiniBand network is now installed and available for operation.
4. **This procedure ends here.**

Installing or replacing a GX Host Channel Adapter

This procedure guides you through the process for installing or replacing an GX Host Channel Adapter (HCA). The process includes:

- Physically installing or replacing the adapter hardware into your system unit.
- Configuring the LPAR profiles with a new globally unique ID for the new adapter in your switch environment.
- Verifying that the HCA is recognized by the operating system.

Notes:

1. If you are considering deferred maintenance of a GX HCA, review “Deferring replacement of a failing Host Channel Adapter” on page 66.
2. If you replace an HCA, it is possible that the new HCA could be defective in a way that prevents the logical partition from activating. In this case, notification pops up on the controlling HMC. If this occurs, decide if you want to replace the ″new-defective″ HCA immediately, or if you want to defer maintenance and continue activating the logical partition. To defer maintenance and continue activating the partition, you must unassign the HCA in all the partition profiles that contain the HCA using the procedure found in “Recovering from an HCA preventing a logical partition from activating” on page 100.

To install or replace a GX HCA, do the following:

1. If you are performing an adapter replacement, first record information about the adapter being replaced. Important information includes: the logical partitions in which it is used, the GUID index used in each logical partition, and the capacity used in each logical partition. Do the following from the Hardware Management Console that manages the server in which the HCA is installed.
   a. Obtain the list of partition profiles that use the HCA. If there is no list, proceed to the next step.
   b. Obtain or record the GUID index and capability settings in the partition profiles that use the HCA:
      1) Go to the **Server and Partition** window.
      2) Select the **Server Management** partition.
      3) Expand the server in which the HCA is installed.
      4) Expand the partitions under the server.

5) Expand each partition that uses the HCA. If you do not know which partition uses the HCA, you must expand the following for each partition profile, and record which ones use the HCA, as well as the GUID index and capability settings.
   a) Select each partition profile that uses the HCA.
   b) From the menu, click **Selected** → **Properties**.
   c) In the **Properties** dialog, click the **HCA** tab.
   d) Using its physical location, find the HCA of interest.
   e) Record the GUID index and capability settings.

2. Install or replace the adapter in the system unit. For instructions on installing a GX HCA in your system unit, see the *RIO/HSL or InfiniBand (IB) adapter* topic in the IBM Systems Hardware Information Center.

   Note: When a Host Channel Adapter (HCA) is added to a logical partition, the HCA becomes a required resource for the partition. If the HCA ever fails in such a way that the system's GARD function prevents it from being used, the logical partition cannot be reactivated. If this occurs, a pop-up message displays on the controlling HMC which indicates that you need to unassign the HCA from the logical partition to continue activation. The GARD function is invoked for serious adapter or bus failures that could impair system operation, such as ECC errors or state machine errors. InfiniBand link errors should not invoke the GARD function.

3. Update the LPAR profiles (for all partitions that will use the new GX HCA) with the new Globally Unique ID (GUID) for the new InfiniBand GX HCA.

   Each GX HCA has a Globally Unique ID (GUID) that is assigned by the manufacturer. If any of these adapters are replaced or moved, the LPAR profiles for all partitions that will use the new GX HCA must be updated with the new GUID. The customer can do this from the HMC that is used to manage the server in which the HCA is installed. Do the following:
   a. Go to the **Server and Partition** window.
   b. Select the **Server Management** partition.
   c. Expand the server in which the HCA is populated.
   d. Expand the Partitions under the server.
   e. Expand each partition that uses the HCA, and perform the following for each partition profile that uses the HCA:
      1) Select each partition profile that uses the HCA.
      2) From the menu, click **Selected** → **Properties**.
      3) In the **Properties** dialog, click the **HCA** tab.
      4) Using its physical location, find and select the HCA of interest.
      5) Click **Configure**.
      6) Enter the GUID index and Capability settings. If this is a new installation, obtain these settings from the installation plan information. If this is a repair, refer to the setting that you previously recorded in step 2.
      7) If the replacement HCA is in a different location than the original HCA, you should now clear out the original HCA information from the partition profile, by choosing the original HCA by its physical location and clicking **Clear**.

   Note: If the following message occurs when you attempt to assign a new unique GUID, you might be able to recover from this error without the help of a service representative.

   ```
   A hardware error has been detected for the
   adapter U787B.001.DNW45FD-P1-Cx.
   You cannot configure the device at this time.
   Contact your service provider
   ```

The Service Focal Point, can be accessed on your HMC, see the ″Start of call″ procedure in the IBM Systems Hardware Information Center, and perform the indicated procedures. Check the Service Focal Point and look for reports that are related to this error. Perform any recovery actions that are indicated. If you cannot recover from this error, contact your service representative.

4. After the server is booted, verify that the HCA is recognized by the operating system. See “Verifying the installed InfiniBand network (fabric) in AIX or Linux” on page 67.
5. You have finished installing and configuring the adapter. If you were directed here from another procedure, return to that procedure.
6. **This procedure ends here.**

Deferring replacement of a failing Host Channel Adapter

If you plan to defer maintenance of a failing Host Channel Adapter (HCA), determine the risk of the HCA failing in such a way that it could prevent future logical partition reactivation. To assess the risk, determine if there is a possibility of the HCA preventing the reactivation of the logical partition. If this is possible, you must consider the probability of a reboot during the time that maintenance is deferred. To determine the risk, do the following:

1. Go to the **Server and Partition** window.
2. Click the **Server Management** partition.
3. Expand the server in which the HCA is installed.
4. Expand the partitions under the server.
5. Expand each partition that uses the HCA. If you do not know which partitions use the HCA, you must expand the following for each partition profile, and record which partitions use the HCA.
   a. Select each partition profile that uses the HCA.
   b. From the menu, click **Selected → Properties**.
   c. In the **Properties** dialog, click the HCA tab.
   d. Using its physical location, locate the HCA of interest.
   e. Verify that the HCA is managed by the HMC.
6. To determine whether to defer maintenance, there are two possibilities:
   - If you find that the HCA is not managed by the HMC, it has failed in such a way that it will be GARDed off during the next IPL. Therefore, consider that until maintenance is performed, any of the partitions using the failed HCA might not properly activate until the HCA is unassigned. This affects future IPLs that the customer wishes to perform during the deferred maintenance period. Also, any other failure that requires a reboot also results in the partition not activating properly. To unassign an HCA, please see “Recovering from an HCA preventing a logical partition from activating” on page 100. If you unassign the adapter while the partition is active, the HCA is actually unassigned at the next reboot.
   - If you find that the HCA is managed by the HMC, the HCA failure will not result in the GARDing of the HCA, and deferred maintenance will not risk the prevention of partition activation because of a GARDed HCA.

Installing a GX Dual-Port Host Channel Adapter in a model 52A, 550, 550Q, 560Q, or 570

For information about installing a model 52A, 550, 550Q, 560Q, or 570 InfiniBand adapter, see the installation instructions in “Installing or replacing a GX Host Channel Adapter” on page 64.

Installing a GX Dual-Port Host Channel adapter in a model 575, 590, or 595

Installing a GX Dual Port Host Channel Adapter (HCA) in a system.

Note: Installing a Dual Port HCA in the 575, 590, or 595 servers is a procedure performed by an authorized service provider.

For procedures about installing an InfiniBand GX Dual Port HCA in a model 575, 590, or 595, see http://w3.rchland.ibm.com/projects/WCII.

Note: This address is intended for use by authorized service representatives.

The HCAs for the 590 or 595 are 12x devices; see the section on "Verifying static-12x mode connectivity."

Verifying the installed InfiniBand network (fabric) in AIX or Linux

Verifying the installed InfiniBand network (fabric) in AIX or Linux

After the InfiniBand network is installed, the GX adapters and the network fabric must be verified through the operating system. To verify the installed InfiniBand network (fabric) in AIX or Linux, refer to the following topics:

- AIX, see "Verifying the GX HCA connectivity in AIX."
- Linux, see "Verifying the GX HCA to InfiniBand fabric connectivity in Linux"

Verifying the GX HCA connectivity in AIX

To verify the GX HCA connectivity in AIX, check the HCA status by running the `lsdev -C | grep ib` script:

An example of good results for verifying a GX HCA is: `iba0 Available GX Host Channel Adapter`.

Verifying the GX HCA to InfiniBand fabric connectivity in Linux

Refer to the documentation provided with SUSE Linux Enterprise Server 9 (SP2) with IBM IB GX HCA driver and Open Fabrics Enterprise Distribution. Refer to the instructions in the eHCAD install file contained within the download from SourceForge web site (http://sourceforge.net/projects/ibmehcad).

Verifying static-12x mode connectivity

This section describes the process for verifying static-12x mode connectivity and operation.

You should have been referred here by the procedure in Verify the InfiniBand network topology and operation.

This procedure describes how to verify static-12x mode connectivity for octopus cable connections on a link configured for static-12x operation.

Verifying static-12x mode connectivity involves verifying LED states and checking port status in the IBM Network Manager's Switch Topology View.

Note: Before continuing with any verification procedures, ensure that everything in your network is configured, connected, and powered-on.

Verifying a 7048-120 or SFS7000P

For a 7048-120 or SFS7000P, if the group is properly configured and the cable connectors properly plugged, all three LEDs will come on with only one that is blinking. Also, only one switch port in the

group should be active in the Switch Topology View. Check the LEDs when the attached server is powered on to at least partition standby state, and the cables are plugged. Then check the IBM Network Manager's Switch Topology View.

The following two tables show the expected LED and IBM Network Manager's Switch Topology View Port Status states for a properly configured static-12x grouping. You must verify all three states because it is possible that one may indicate success and another indicate failure. The first table is for groups where the lowest numbered port is connected to 16, which indicates that the lowest numbered port should be the active port. The second table is for groups where the highest numbered port is connected to 16, which indicates that the highest numbered port should be the active port.

Table 17. Group ports LED status when the lowest numbered switch port in the group should have connector 16 attached

	Lowest Numbered Port	**Middle Port**	**Highest numbered Port**
Connectors	16	17	18
Port LED (green)	On blinking	On	On
IBM Network Manager	Active	Down	Down

Note: If you shift the group of three 4x connectors to the left or to the right one or more ports, it is possible to establish a connection on one or more of the ports adjacent to the group, if the adjacent group of ports is not configured to static-12x mode. However, you will not see all three green LEDs come on unless all 3 connectors are outside of the intended group.

Table 18. Group ports LED status when the lowest numbered switch port in the group should have connector 18 attached

	Lowest Numbered Port	**Middle Port**	**Highest numbered Port**
Connectors	18	17	16
Port LED (green)	On	On	On blinking
IBM Network Manager	Down	Down	Active

Note: If you shift the group of three 4x connectors to the left or to the right one or more ports, it is possible to establish a connection on one or more of the ports adjacent to the group, if the adjacent group of ports is not configured to static-12x mode. However, you will not see all three green LEDs come on unless all 3 connectors are outside of the intended group.

Problems Found in Verifying a 7048-120 or SFS7000P

If connection problems occur with a switch, check the following:

1. Verify with the CLI that the lowest numbered port in the group is configured to: `admin-speed=12x and auto-negotiate=disable`.
2. Verify that the 4x connectors are plugged into the correct ports. See the cable plan, or "Planning octopus cables in static 12x cabling" on page 19 and "Planning 4x connections for octopus cables" on page 21. Ensure that the cables have not been shifted left or right one or two ports, and that the order has not been reversed or jumbled.
3. Reseat the cable connectors.
4. If only one group is having difficulty, the cable may be bad. Try another cable.
5. If you have been configuring several groups and plugging (and maybe unplugging) cable connectors on a network that has been up and running, consider the following:
   a. If the only issue is that the IBM Network Manager does not show Active port status, you may wish to restart the IBM Network Manager.
   b. Otherwise, you may have to reboot the switch and re-power the servers.

1) Refer to “Rebooting the entire switch chassis” on page 145. If you feel you must re-power the switch, you must pull and re-plug all power cables for the switch.
2) You can re-power the servers to partition-standby or automatically start the partitions.

6. Replace the switch, or the adapter.
7. After any repairs to a link, go back to the beginning of “Verifying static-12x mode connectivity” on page 67.
8. Return to the detailed procedure in “Verify the InfiniBand network topology and operation” on page 62

Verifying a 7048-270 or SFS7008P

If the group is properly configured and the cable connectors properly plugged, all three Traffic LEDs will come on, and only one Logical/Link LED should be on or blinking. Furthermore, only one switch port in the group should be Active in the Switch Topology View. Check the LEDs when the attached server is powered on to at least partition standby, and the cables are connected. Then, check the IBM Network Manager’s Switch Topology View.

The following two tables show the expected LED and IBM NM Switch Topology View Port Status states for a properly configured static-12x grouping. It is imperative that all three states be verified, because it is possible that one may indicate success and another indicate failure. The first table is for groups where the lowest numbered port is connected to 16, which indicates that the lowest numbered port should be the active port. The second table is for groups where the highest numbered port is connected to 16, which indicates that the highest numbered port should be the active port.

Table 19. Group ports LED status when the lowest numbered switch port in the group should have connector 16 attached

	Lowest Numbered Port	**Middle Port**	**Highest numbered Port**
Connectors	16	17	18
Traffic LED (green)	On blinking	On	On
Logical LED (green)	On	Off	Off
IBM Network Manager	Active	Down	Down

Note:

1. If the traffic LED is On, see “Interpreting LEDs” on page 122.
2. If you shift the group of three 4x connectors to the left or to the right one or more ports, it is possible to establish a connection on one or more of the ports adjacent to the group, if the adjacent group of ports is not configured to static-12x mode. However, you will not see all three green LEDs come on unless all three connectors are outside of the intended group.

Note:

Table 20. Group ports LED status when the lowest numbered switch port in the group should have connector 18 attached

	Lowest Numbered Port	**Middle Port**	**Highest numbered Port**
Connectors	18	17	16
Traffic LED (green)	On	On	On blinking
Logical LED (green)	Off	Off	On
IBM Network Manager	Down	Down	Active

Note:

1. If the traffic LED is On, see “Interpreting LEDs” on page 122.
2. If you shift the group of three 4x connectors to the left or to the right one or more ports, it is possible to establish a connection on one or more of the ports adjacent to the group, if the adjacent group of ports is not configured to static-12x mode. However, you will not see all three green LEDs come on unless all 3 connectors are outside of the intended group.

Problems Found when verifying a 7048-270 or SFS7008P

If connection problems occur with a switch, check the following:

1. Verify with the CLI that the lowest numbered port in the group is configured to: `admin-speed=12x and auto-negotiate=disable`
2. Verify that the 4x connectors are plugged into the correct ports. See “Planning 4x connections for octopus cables” on page 21. Make sure that the cables have not been shifted left or right one or two ports, and that the order has not been reversed or jumbled.
3. Reseat the cable connectors.
4. If only one group is having difficulty, the cable may be bad. Try another cable.
5. If you have been configuring several groups and plugging (and maybe unplugging) cable connectors on a network that has been up and running, consider the following:
   a. If the only issue is that the IBM Network Manager does not show active port status, restart the IBM Network Manager.
   b. Otherwise, you may have to reboot the switch and re-power the servers.
      1) Refer to the “Rebooting the entire switch chassis” on page 145. If you feel you must re-power the switch, you must pull and re-plug all power cables for the switch.
      2) You can re-power the servers to partition-standby or automatically start the partitions.
6. Replace the LIM.
7. Replace the adapter. Go to “Installing or replacing a GX Host Channel Adapter” on page 64

   Note: If you are considering deferred maintenance of the adapter, review “Deferring replacement of a failing Host Channel Adapter” on page 66.

If you came here from “Verifying the installed InfiniBand network (fabric) in AIX or Linux” on page 67, then return to that section now. Otherwise, return to the beginning of “Verifying static-12x mode connectivity” on page 67 and verify that the repair has rendered the link operational.

Runtime errors

Report run time errors are reported to Service Focal Point with the appropriate FRU lists. You should refer to the verification section after any repairs to a link.

Managing InfiniBand networks with IBM Network Manager

Cluster networks that are connected with Host Channel Adapters (HCAs) and InfiniBand switches can be managed through the IBM Network Manager, and the Element Manager available from TopSpin. However, any PCI HCAs on your network are not available to the IBM Network Manager.

Managing GX HCA-based InfiniBand networks

Use the IBM Network Manager to manage your InfiniBand network. The IBM Network Manager enables you to manage your InfiniBand network from the Hardware Management Console (HMC).

Use the IBM Network Manager to manage your InfiniBand switches, update switch software, view network topology information, and view and modify management properties. You can also view the Network Manager event logs. You must enable the IBM Network Manager from the HMC before you can use it to manage your network.

The remainder of this section covers the various windows in the IBM Network Manager. For detailed information on how to interpret the status seen in these windows, see "Status procedures for the IBM Network Manager" on page 189.

Enabling and disabling the IBM Network Manager

Enable the IBM Network Manager to manage your InfiniBand switch network.

Before you can use the IBM Network Manager to manage your InfiniBand network switches and servers, you must enable it from the HMC. When the IBM Network Manager is enabled, it starts a discovery process on the network and begins providing data about the status of your InfiniBand switch network.

Any user can enable or disable the IBM Network Manager.

Note: Only one instance of IBM Network Manager can be run in a cluster. Unpredictable results will occur if more than one instance is running.

If you decide to move the IBM Network Manager functions from one HMC host to another HMC host, you must first disable the IBM Network Manager on the first HMC host before enabling it on the second host.

To enable or disable the IBM Network Manager, complete the following steps:

1. In the Navigation area, expand the **Switch Management** folder.
2. Click **IBM Network Manager**.
3. In the Network Manager menu, select **Enable IBM NM Software**.
   To disable the IBM Network Manager, select **Disable IBM NM Software**.

Viewing switch topology information in an InfiniBand network

View information about the switch layout and connectivity in your InfiniBand network.

You can view information about the switches in your InfiniBand network that identifies the physical location and connection activity for all connected switch devices.

Any user can view the switch topology information.

© Copyright IBM Corp. 2004, 2007

Note: To access this task on the HMC, you must first enable the IBM Network Manager.

To view the switch topology information, complete the following steps:

1. In the Navigation area, expand the **Switch Management** folder.
2. Click **IBM Network Manager**.
3. From the Network Manager menu, select **View Switch Topology**.

   Note: It is possible for the menu option **Selected-Expand/Collapse** and the twisty indicator in the window to get out of sync, where one may act as if the selected object is expanded, while the other may act as if the selected object is collapsed. If this happens, select another object, then reselect the object of interest.

   The HMC displays a table with a list of the switches in the InfiniBand network. By default, the information presented shows the switch device name (or identifier), location code, service subsystem connectivity status, GID-prefix, GUID, port status, and neighbor port (link) information.

   Note: The physical location information for any PCI HCA is not displayed in the switch topology information.

   On the bottom left of the window you will see the timestamp for the last auto-update. This represents the last time at which IBM Network Manager updated the information in the Switch Topology view.

   To change the columns displayed in the list of information, click **View → Show Columns** and select or clear the check box for the columns that you want to designate as the default columns.

Viewing server topology information in an InfiniBand network

View information about the physical layout of your InfiniBand network.

Any user can view the server topology information.

You can view information about the servers, adapters, and ports in your InfiniBand network. Currently the IBM Network Manager only supports GX host channel adapters in this view. No information will be available for PCI host channel adapters and these adapters will not be displayed.

Note: To access this task on the HMC, you must first enable the IBM Network Manager.

To view the server topology information, complete the following steps:

1. In the Navigation area, expand the **Switch Management** folder.
2. Click **IBM Network Manager**.
3. From the Network Manager menu, select **View End-Point Topology**.

   Note: It is possible for the menu option **Selected-Expand/Collapse** and the twistie indicator in the window to get out of synch, where one may act as if the selected object is expanded, while the other may act as if the selected object is collapsed. If this happens, simply select another object, then reselect the object of interest.

   The HMC displays a table with a list of the servers in the InfiniBand cluster. By default, the information presented shows the server name, location code, power on/off status, adapter status, port status, and neighbor port (link) information.

   Note:

   a. Servers will be displayed with their machine type model and serial (MTMS) in the location code column, however, HCAs and their ports will have IBM location codes in the location code column.
   b. You will see servers with PCI HCAs. Because the IBM Network Manager does not gather information about PCI HCAs, you cannot expand down to the PCI HCAs.

On the bottom left of the window you will see the timestamp for the last auto-update. This represents the last time at which IBM Network Manager updated the information in the End-Point view.

To change the columns displayed in the list of information, click **View → Show Columns** and select or clear the check box for the columns that you want to designate as the default columns.

Viewing logical topology information in an InfiniBand network

View information about the logical server layout in your InfiniBand network.

You can view information about the servers, adapters, ports, and related status for the logical partitions in your InfiniBand network.

Any user can view the logical topology information.

Note:

1. To access this procedure, you must first enable the IBM Network Manager.
2. Servers are displayed with their machine type model and serial (MTMS) in the location code column, however, HCAs and their ports have IBM location codes in the location code column.

To view the logical topology information, complete the following steps:

1. In the Navigation area, expand the **Switch Management** folder.
2. Click **IBM Network Manager**.
3. From the IBM Network Manager menu, select **View Logical Topology**.

   Note: It is possible for the menu option **Selected-Expand/Collapse** and the twisty indicator in the window to get out of sync, where one may act as if the selected object is expanded, while the other may act as if the selected object is collapsed. If this happens, select another object, then reselect the object of interest.

   The HMC displays a table with a list of the servers and partitions in the InfiniBand network configuration. By default, the information presented shows the server name, location code, system or logical partition status, adapter status, adapter type, the GUID, port status, and neighbor port (link) information.

   On the bottom left of the window you will see the timestamp for the last auto-update. This represents the last time at which IBM Network Manager updated the information in the Logical Topology view.

   To change the columns displayed in the list of information, click **View → Show Columns** and select or clear the check box for the columns that you want to designate as the default columns.

   Note: You will see only GX HCAs in this display.

   The physical GX host channel adapter comprises logical HCAs and logical switches on a physical adapter. This structure allows InfiniBand switch partitioning of logical HCAs. When viewing a server and its physical HCAs through the Logical Topology view, the logical switches exist at all times and provide the connectivity from logical HCAs to the physical switch network. The first port of a logical switch connects the server to the physical switch network.

   The following figure shows an example of how logical switches can be configured to enable logical partitions to connect to the physical switch network.

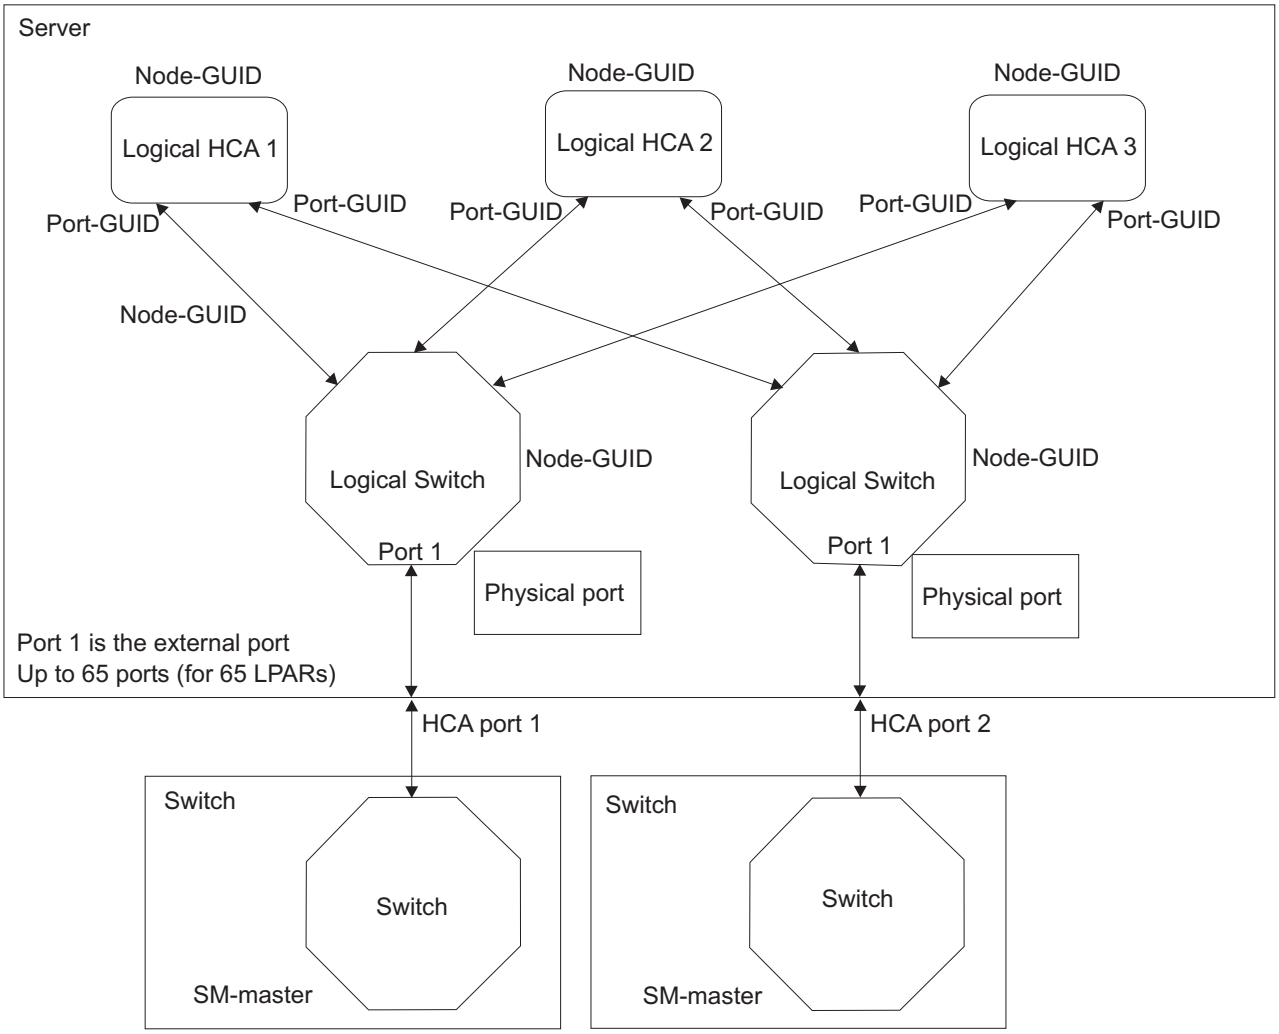

Figure 4. Connection of logical partitions to logical switches on a GX host channel adapter.

Until a logical HCA exists on a physical HCA, you will see only the physical HCA and the first port, which connects to the physical switch network, of each logical switch.

Logical HCAs exist only after a physical HCA has been assigned to at least one logical partition. The logical HCA does not appear in the Logical Topology View until the logical HCA exists. Also, because the second port of a logical switch provides the connection to the logical HCA. Before a logical HCA exists on a physical HCA, the Logical Topology View will not display port number 2 of each logical switch.

Because of the independence between the first and second port of the logical switch, it is possible for the connection between a logical HCA and a logical switch to have a status that is different from the logical switch port number 1 (which connects to the physical switch network).

Viewing IBM Network Manager properties

View and manage switch information and the properties of the currently enabled IBM Network Manager for your InfiniBand environment.

You can view and modify the IBM Network Manager properties and the switch-management properties for your InfiniBand network. For example, use this procedure to change the IBM Network Manager default name assigned to the managed switches. Using switch names provides a convenient way to keep track of the switches that you are managing, particularly when frame and unit location codes are not

readily available. You can also change the switch-management priority (change the subnet manager master and succession order), synchronize the switch local time with the HMC time, and view the switch topology.

Any user can view the Network Manager properties.

Note: The Network Manager must be enabled in order for you to access this task from the HMC.

To view the properties, complete the following steps:

1. In the Navigation area, expand the **Switch Management** folder.
2. Click **IBM Network Manager**.
3. From the Network Manager menu, select **View Management Properties**.
4. Click the **IBM NM** tab to view the IBM Network Manager properties.
   The HMC displays the version of IBM Network Manager (IBM-NM) and the IBM Network Manager Problem Analysis plugin (IBM-NM PA) for the local HMC.
5. Click **Switch** to view and modify information about the switches in your InfiniBand environment. On the bottom left of the window you will see the timestamp for the last auto-update. This represents the last time at which IBM Network Manager updated the information in the Management Properties view.

Viewing the IBM Network Manager event log

View the IBM Network Manager event log.

The IBM Network Manager maintains an event log that you can view to track the IBM Network Manager activities.

Any user can view the event logs.

Note: To access this task in the HMC the IBM Network Manager must be enabled.

To view the IBM Network Manager event log, complete the following steps:

1. In the Navigation area, expand the **Switch Management** folder.
2. Click **IBM Network Manager**.
3. From the Network Manager menu, select **View IBM NM event log**.
   The HMC displays the contents of the log file.

Updating switch software

Use the IBM Network Manager to update the software for all the switches in your network.

If your Cisco or Topspin switches are not at a firmware level of 2.5.0 or higher, you must first enable ftp (file transfer protocol) on the HMC on which you are working. Go to “Updating switch software from 2.3.0 or earlier to 2.5.0 or later” on page 76.

You can update multiple switches simultaneously for a single software version, or choose a switch and the version of software to apply for the update of that switch. You can start another switch installation before a switch update procedure that is currently running has finished. However, if both installations run concurrently and contain the same switch, the subsequent installation action might fail.

An Import option allows you to add more software versions from a DVD.

Any user can update the switch software.

Note: To access this task from the HMC, you must first enable the IBM Network Manager.

To update your switch software, complete the following steps:

1. In the Navigation area, expand the **Switch Management** folder.
2. Click **IBM Network Manager**.
3. From the IBM Network Manager menu, select **Update Switch Software**.
4. Select the switches that you want to update.
5. Select the software version to install for the selected switch or switches.
6. Click **OK** to start the software update.

Note: To uninstall the last update, select **Return switch to previous software**.

Updating switch software from 2.3.0 or earlier to 2.5.0 or later

Before attempting to upgrade your switch firmware with the IBM Network Manager, you need to upgrade Cisco or Topspin switch firmware from version 2.3.0 or earlier to version 2.5.0 or later.

FTP must be enabled on the HMC running IBM Network Manager (IBM NM) to upgrade switch firmware from version 2.3.0 or earlier.

If you are need to upgrade Cisco or Topspin switch firmware from version 2.3.0 or earlier to version 2.5.0 or later, use the following procedure:

1. Enable ftp (file transfer protocol) on the HMC by:
   a. Obtain root access to the HMC. This may require you to log in using PESH:
      1) From the HMC desktop, right click and select **Terminals** and then **rshterm**.
      2) From the command prompt, enter `lshmc -v` (this command returns information about the HMC on which you are working). You need the serial number of the HMC. The data returned will contain a field labeled, **SE**. This is the serial number of the HMC. Record this information.
      3) Contact Support for the 24 hour PE passcode. Provide Support the current date on HMC and the serial number of the HMC. The date/time reference for the passcode is based on GMT (Greenwich Mean Time).
      4) On the HMC, enter `pesh` *`hmcserialnumber`*.
      5) When prompted, enter the passcode that was generated from the Support/PE.
   b. Use the vi editor to edit gssftp. Enter `vi /etc/xinetd.d/gssftp`
   c. Change `disable = yes` to `disable = no`.
   d. Save the file by entering `wq` on the vi command line.
   e. To enable ftp, enter `/etc/init.d/xinetd restart` from the command prompt.
2. Retry the switch firmware update from IBM Network Manager, Management Properties View panel. Go to "Updating switch software" on page 75. After updating the switches firmware, disable the ftp by continuing with the next step in this procedure.
3. To disable ftp on the HMC do the following:
   a. Obtain root access to the HMC.
   b. Use the vi editor to edit gssftp. Enter `vi /etc/xinetd.d/gssftp`
   c. Change `disable = no` to `disable = yes`.
   d. Save the file by entering `wq` on the vi command line.
   e. To disable ftp (file transfer protocol), enter `/etc/init.d/xinetd restart` from the command prompt.

InfiniBand network problem determination

Tools are available to help troubleshoot problems on an InfiniBand network. The primary tool is the IBM Network Manager..

Note: IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.

To begin troubleshooting an InfiniBand network problem, do the following:

1. Are you troubleshooting a problem on a network that is managed by the IBM Network Manager?

   No: Go to step 2 on page 78

   Yes: The following table contains an index of symptoms and procedures to follow for each symptom. Read through the symptoms and do the appropriate action. If you cannot determine the source of the problem, go to 2 on page 78

Symptom	Action
CBxxxxxx 8-digit reference code. See the note following this table.	Record the isolation procedure that is listed as the first FRU in the FRU list for the serviceable event, and refer to "InfiniBand problem isolation procedures" on page 101. If there are multiple links down on the same card or cards that are together in a LIM pair (see "Plugging Rules for 7048-270 or SFS7008P switch line interface modules and fabric controllers" on page 152, use the procedure found in "Determining faulty fabric controller cards versus faulty LIM cards" on page 153.If no isolation procedure was given, call your next level of support.
All other reference codes	Use the Repair and Verify procedures on the HMC, or refer to the IBM Systems Hardware Information Center's Start of call procedure.
Status in IBM Network Manager interface.	Refer to "Status procedures for the IBM Network Manager" on page 189.
Switch will not boot	Refer to "Diagnosing an InfiniBand switch that will not boot" on page 101.
Switch LED indicating problem	Refer to "Interpreting LEDs" on page 122.
Device powering off	Refer to "Suspected power problem" on page 139.
Device appears hot	Refer to "Suspected thermal problem" on page 140.
A logical partition does not activate and a notification pops up on the controlling HMC which indicates that a Host Channel Adapter (HCA) is not operational.	You must unassign the HCA to continue activation of the logical partition. Use the procedure found in "Recovering from an HCA preventing a logical partition from activating" on page 100.
Suspect a specific link is bad	Refer to "Isolating a problem with a port or link" on page 136.
Performance problems	Refer to "Isolating performance problems" on page 138.
LIMs for a 7048-270 or SFS7008P are missing from the Switch Topology View	Refer to "Missing LIM(s) or LIM ports" on page 153.

© Copyright IBM Corp. 2004, 2007

Symptom	Action
Many devices are not found in the IBM Network Manager Views, or are no longer Responsive, or have Unknown status values.	1. Check the service network connections for the HMC and the devices. 2. Check any Ethernet switches or routers used for the service network.
The GID-prefix for a switch is incorrect.	Use the procedure in “Switch is on the incorrect subnet” on page 144
The LMC is incorrect for a switch	Use the procedure in “Setting the location identifier mask control” on page 148.
Inconsistent Results from IBM Network Manager (like disappearing devices)	If the problem does not appear to be the service network, verify that IBM Network Manager is running on only a single HMC. It is possible to enable it on multiple HMCs, in which case, the different instances of IBM Network Manager are attempting to manage the network.
Disappearing IBM Network Manager windows	“Disappearing IBM Network Manager windows” on page 152
Cannot Expand or Collapse an object using the menu option **Selected-Expand/Collapse** or the indicator in the window.	Select another object and reselect this object.

Note:

- Reference codes based on serviceable events detected by the InfiniBand switches (CB*xxxxxxx* reference codes) do not support the Call Home feature and these events will not be called home by the Electronic Service Agent. If you attempt to use the manual function to initiate Call Home, it will not successfully generate a service call, and you will not receive any indication that the service call was not generated. However, if necessary, it is possible to manually Call Home the extended error data.
- After repair procedures, verify the network function according to the “Repair verification” on page 150 procedure.

2. The following table contains an index of symptoms and procedures to follow for each symptom. Is the problem on the network described by any of the symptoms in the table?

   No: call your next level of support.

   Yes: The following table contains an index of symptoms and procedures to follow for each symptom. Read through the symptoms and do the appropriate action. If you cannot determine the source of the problem, call your next level of support.

Symptom	Procedure
All reference codes that are found in the Service Focal Point.	This situation applies only to a cluster with an HMC. Use Repair and Verify procedures, or refer to the IBM Systems Hardware Information Center’s Start-of-call procedure or reference codes section.
All other reference codes, SRNs, or SRCs	Refer to the IBM Systems Hardware Information Center’s Start-of-call procedure or reference codes section.
InfiniBand Error in OS error log	Refer to the IBM Systems Hardware Information Center’s Start-of-call procedure.
Switch will not boot	1. Check the power. 2. Check the Element Manager for an indication of a switch failure or error status. 3. Check the ts_log.

Symptom	Procedure
An LED indicates problem	Refer to the hardware documentation for the type of switch that has the LED that is indicating a problem.
Device powering off	1. Check the LEDs on the power supplies. 2. Check the power cable connections. 3. Check the Element Manager status of the power supplies. 4. Verify that the site meets power requirements.
Device appears hot	1. Check the LEDs on the fans. 2. Verify that the fans are rotating. 3. Check the Element Manager status of the fans and sensors. 4. Verify that the site meets cooling specifications.
Suspect a specific link is bad	1. Check the port statistics counters on the link. 2. If any errors are non-zero, watch them and determine if they are increasing. If they are, the link is bad. 3. If errors, check for loose cables or damaged pins. 4. Replace FRUs in the following order: a. cables b. HCA
Status in Element Manager	See the *Element Manager User Guide*, order number: 10-00116-02-A0.
Performance problems	1. Look for hardware problems: a. Check the error logs on the servers. b. Check the Element Manager status. c. Check the LEDs on the switches and HCAs. d. Check the ts_log on the Element Manager. 2. Check for missing resources. 3. Verify that all processors and memory are configured on the servers. 4. Check for configuration problems: a. Check the configuration in the Element Manager and verify that the network is cabled as expected. b. Verify that the PCIx HCAs are visible to the operating system.

Note: Reference codes based on serviceable events detected by the InfiniBand switches (CB*xxxxxxx* reference codes) do not support the Call Home feature and these events will not be called home by the Electronic Service Agent. If you attempt to use the manual function to initiate Call Home, it will not successfully generate a service call, and you will not receive any indication that the service call was not generated. However, if necessary, it is possible to manually Call Home the extended error data.

InfiniBand fabric maintenance strategy

This section contains links to service and support topics for setting up a clustered environment and getting fixes for a clustered server.

The following table contains links to topics instructing you how to set up connectivity to service and support and how to get fixes for the clustered server.

Maintenance task	Procedure
Learn about all the tasks you need to perform to set up connectivity to service and support when you are working in a clustered environment.	"Setting up a clustered environment to connect to service and support"
Learn how to get fixes for the clustered server.	"Getting fixes for a clustered environment" on page 91

Setting up a clustered environment to connect to service and support

Learn about all the tasks you need to perform to set up connectivity to service and support when you are working in a clustered environment.

Task 1. Before you begin

This procedure contains the complete list of tasks needed to set up connectivity to service and support. Some of these tasks might already have been completed (during initial server setup, for example). If so, you can use this procedure to verify that the tasks were completed correctly.

In this document, a *direct Internet connection* is defined as access to the Internet from a logical partition, server, or HMC by direct or indirect access. *Indirect* means that you are behind a Network Address Translation (NAT) firewall. *Direct* means that you have a globally routable address without an intervening firewall (which would block the ports that are needed for communication to service and support).

Task 2. Determine your connectivity method

Choose the method that best describes your situation.

If you have an HMC and have multiple logical partitions:

- For the HMC, use either a direct Internet or dial-up connection to connect the HMC to service and support.
- For AIX or Linux logical partitions, hardware errors will be reported through the HMC, using the connection method provided for the HMC.

If you do not have an HMC but you have AIX or Linux:

- If you do not have logical partitions, use a direct Internet, direct dial-up connection, or Secure Sockets Layer (SSL).
- If you have logical partitions and are using the Integrated Virtualization Manager to manage your server, you might want to configure the service processor to contact service and support when the server is not available. For details, refer to "Task 14. Configure the service processor" on page 89.

  Note: If you are using the Integrated Virtualization Manager to manage your server, you will need to check Service Focal Point for the Integrated Virtualization Manager to see if you need to contact service and support.

© Copyright IBM Corp. 2004, 2007

Task 3. Prerequisites

1. For direct Internet connections, work with the network administrator to verify the following:
   - For HMC environments ensure that the following ports are open for communication:
     - Protocol UDP ports 500 and 4500 with the following IP addresses: Boulder: `207.25.252.196` and Rochester: `129.42.160.16`
     - ESP (protocol 50) with the following IP addresses: Boulder: `207.25.252.196` and Rochester: `129.42.160.16`
   - For IBM eServer p5 or System p5 servers in a non-HMC environment, ensure that the following port is open for communication: Protocol TCP port 443 with the following IP addresses: Boulder: `207.25.252.200` and Rochester: `129.42.160.48`
   - If multiple logical partitions are sharing an Internet connection, you will need the IP addresses or host names created for TCP/IP and for virtual Ethernet.
2. For a dial-up (modem) connection, determine necessary configuration settings, including:
   - Local area code
   - Pre-dial information, such as dialing ″9″ to dial outside the network
   - Use of commas if delayed dialing is needed
3. Ensure that TCP/IP is set up and configured correctly. If not, work with the network administrator and your operating system documentation.

Task 4. Ensure that your physical networking is set up correctly

The underlying framework of your service environment is networking. The following networking connections are required for you to be able to take advantage of electronic services such as reporting hardware problems and other server information and downloading fixes:

- Between the service processor and the HMC
- Between the HMC and the server (AIX and Linux)
- Between your site and service and support

1. Verify the physical connection between the service processor and the HMC.

   The service processor is part of your platform hardware and monitors hardware attributes and conditions on your server. The service processor is controlled by server firmware (Licensed Internal Code) and does not require an operating system to perform its tasks. The connection to the service processor is recommended for all servers, whether or not you have logical partitions. This connection is represented in the following illustration:

Figure 5. This diagram shows the Ethernet connection between your HMC and the service processor on your server.

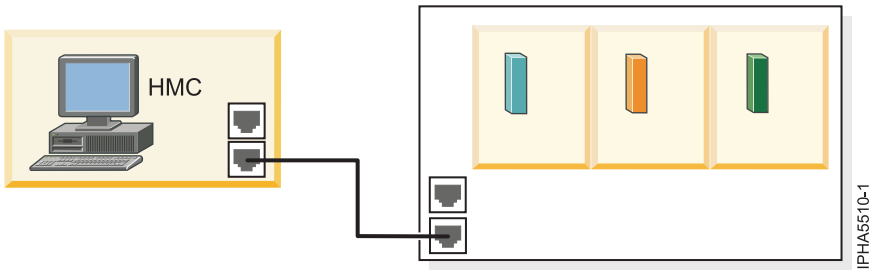

2. Verify the physical connection between the HMC and the server (AIX and Linux).

   This connection allows your server to communicate with your HMC.

   How you set up this connection depends on your configuration:

- If your server is in its manufacturing default configuration, you will make this connection when you set up your server.
- If your server has multiple logical partitions, you must ensure that your HMC can communicate with each logical partition and that the logical partitions can communicate with each other. You will set up these connections as you create your logical partitions.

You can use either of the following methods:

Note: Both of the following networking methods require basic TCP/IP configuration on your logical partitions. For instructions on how to configure TCP/IP, see your operating system documentation.

- Have an Ethernet adapter for one logical partition, most likely your service partition, and then use virtual Ethernet to enable the logical partitions to communicate with each other and with the HMC. This option is the preferred option because it requires that you have only one physical adapter in the system. The following illustration shows this configuration:

Figure 6. This diagram shows the Virtual Ethernet connection between your logical partitions and the physical Ethernet connection between your service partition and the HMC.

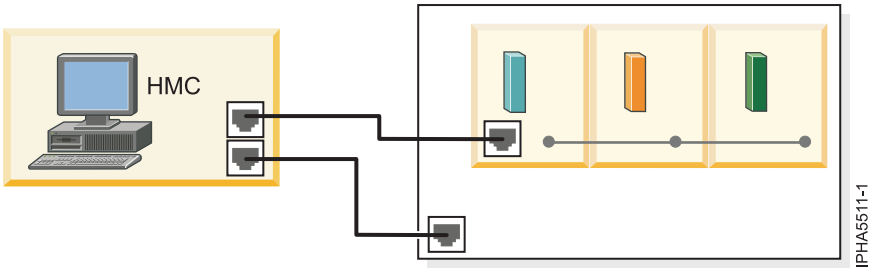

- Have a LAN adapter for each logical partition then have a physical connection between each logical partition and the HMC. This option requires that you have a router and a physical LAN adapter for each logical partition. The following illustration shows this configuration:

Figure 7. This diagram shows physical Ethernet connections between your logical partitions and your HMC through a router.

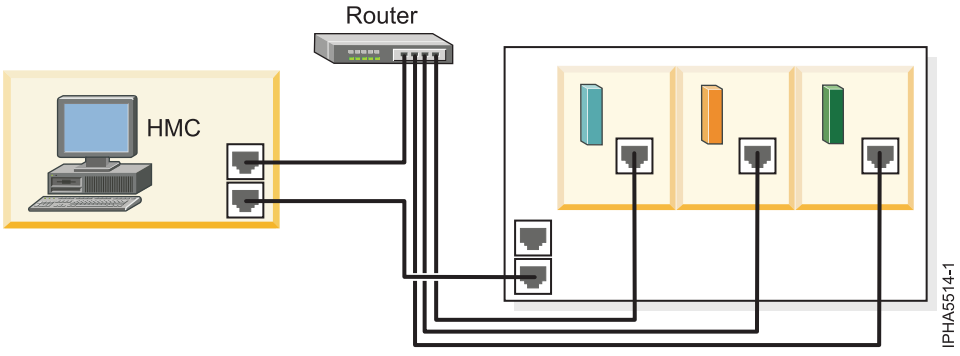

3. Verify the physical connection from your site to service and support.

   This connection enables you to report hardware problems and other server information to service and support. It also enables you to install fixes. This connection is represented in the following illustration:

Figure 8. This diagram shows the connection between service and support and a company that has a server and an HMC.

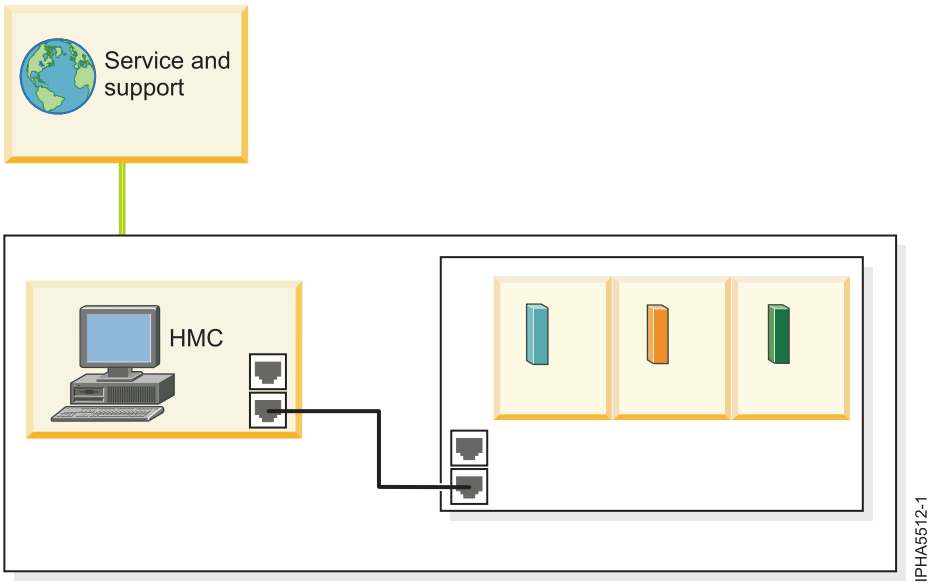

Task 5. Obtain or verify an IBM ID

You will need an IBM ID to register IBM Electronic Service Agent™ on the HMC and for Electronic Service Agent on the operating systems, including AIX, i5/OS™, and Linux. You also will need this ID to view information that has been reported to IBM through Electronic Service Agent.

1. Go to the My IBM Profile (https://www.ibm.com/account/profile) Web site.
2. Verify that you are registered.
   - If you are registered, `Welcome back` will appear on the Web site. Or, you can select **Sign in** and see if your e-mail address is recognized.
   - If you are not registered, select **Register** and fill out the registration form. Create an IBM ID for each of the people you want to have access to the information that Electronic Service Agent reports to IBM. You must associate these accounts with a server, usually your central server. (You can add other servers later.) The people for whom you create IDs must have system administrator authority on all registered servers.
3. Record your IBM ID (the e-mail address that you registered).
4. Go to "Task 6. Verify the HMC service settings using the Guided Setup wizard"where you will use the IBM ID.

Task 6. Verify the HMC service settings using the Guided Setup wizard

The simplest way to verify that the HMC service settings are set up correctly is by using the Guided Setup wizard.

1. Access the Guided Setup wizard using the HMC interface:
   a. In the navigation area, select the HMC that you want to work with.
   b. Click **Information Center and Setup Wizard**.
   c. In the contents pane, click **Launch the Guided Setup Wizard**. The Guided Setup wizard steps you through the tasks that are required to set up your HMC, including the tasks that are required to set up your service environment.
2. Click **Next** to skip the tasks that are not specific to setting up service, including:

- Setting the date and time
- Changing passwords for the hscroot and root user IDs
- Creating user IDs and passwords for new users and setting their authorities
- Specifying network settings

3. Ensure that the following service tasks are completed correctly:
   a. Customer contact information for service-related activities, including:
      - Company name
      - Administrator name
      - E-mail address
      - Phone numbers
      - Information regarding the location of the HMC
   b. Configuration of connectivity for service-related activities.
      - **VPN**

        Note: When configuring the HMC's network settings for connecting using direct or indirect Internet, the HMC must be configured with a default gateway to access the Internet. Select **HMC Management > HMC Configuration > Customize Network Settings**. Ensure that the Default Gateway Information field has a Gateway address listed and a selection is made in the Gateway device field (for example, any).
      - **Dial-up connection from the local HMC**
      - **Connecting through other systems or logical partitions**
   c. Configuration of the network settings.
      - For direct or indirect Internet:
        – HMC host name
        – Domain name
        – Description of HMC
      - For a dial-up modem connection:
        – Dial prefix if applicable
        – Modem configuration, including:
          - Dial type
          - Dial prefix (if applicable)
        – Phone number
   d. Authorize two users for Electronic Service Agent by entering the ID (the e-mail address that you registered with the My Profile Web site, https://www.ibm.com/account/profile).

      Note: You will be able to authorize more users later.
   e. Add e-mail addresses for those you want to be notified when problem events occur.
4. To test the connection from the HMC, open **Service Applications → Remote Support → Customize Outbound Connectivity**.
5. Select the tab for the type of outbound connectivity that you chose for your HMC (Local Modem, Internet VPN, or Pass-Through Systems).
6. Click **Test**.
7. Choose from the following options:
   - If the test is successful, continue with the next Task.
   - If the test fails, continue with the next step.

8. Ensure that your country or region is listed. Select **Service Applications** → **Remote Support** → **Customize Customer Information**. Ensure that your country or region is selected from the drop-down list.
9. Then, choose from the following options:
   - If you have a dial-up connection, perform the following:
     - Check the telephone line going into the HMC and the wall socket.
     - Check to make sure that you have the telephone number configured correctly, including pre-dial information, such as dialing ″9″ to dial outside the network.
   - If you have an Internet VPN connection, perform the following:
     - Ensure that the appropriate firewall rules have been added, if necessary.
     - Check that you have a default gateway set up on the HMC. Select **HMC Management** → **HMC Configuration** → **Customize Network Settings**. Ensure that the Default Gateway Information field has a Gateway address listed and a selection is made in the Gateway device field (for example, any).

Task 7. Set up and configure your logical partitions

For details, refer to the Partitioning the server topic.

Task 8. Install the operating systems on your server or logical partitions

For details, refer to the Installing operating systems topic.

Task 9. Configure your TCP/IP network

For instructions, refer to the operating system documentation.

Task 10. Activate TCP/IP on your server or logical partitions

TCP/IP starts automatically, as long as the network adapter is recognized and can communicate with the network when the AIX or Linux operating system is started.

Task 11. Configure AIX or Linux for connectivity

If you have an HMC, you do not need to configure Electronic Service Agent on AIX or Linux. The AIX and Linux inventory and hardware-problem information (or report, perhaps) are sent through the HMC. However, you might want to set up Electronic Service Agent on AIX or Linux to contact the software service organization.

1. Choose from the following options:
   - If you have an HMC and do *not* want to install Electronic Service Agent, continue with “Task 16. View the server information that was reported to IBM” on page 91.
   - To configure AIX for connectivity, continue with the next step.
   - To configure Linux for connectivity, go to step 8 on page 88.
2. To configure AIX, review the following information:

   In this example, you will configure the following:
   - Electronic Service Agent on the server that has a modem for dial-up connection to service and support

Note: Alternatively, you can use an Internet or Secure Sockets Layer (SSL) connection instead of the modem to connect to service and support.

- Electronic Service Agent on the other clients to communicate with the server that has the modem

Figure 9. This diagram shows three servers and their connection through a modem to service and support.

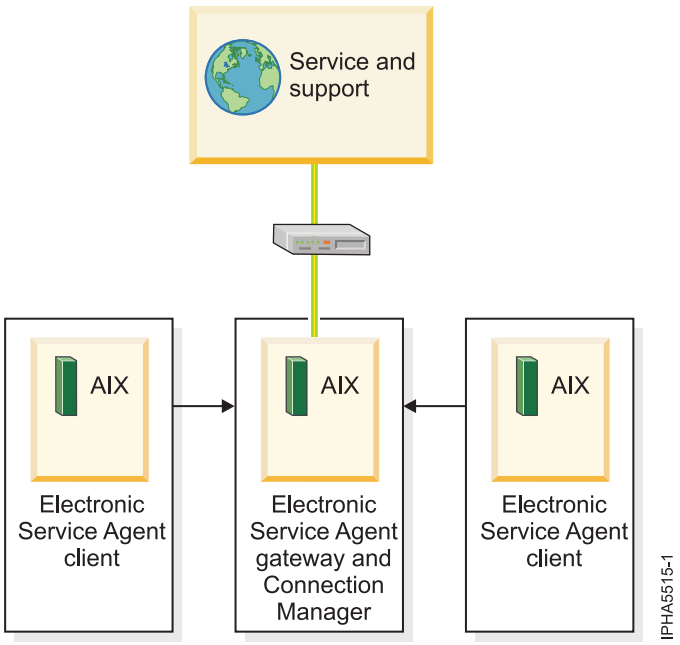

3. Note that for complex network environments involving HTTP proxies, SOCKS proxies, or DMZs, refer to the Electronic Service Agent for IBM pSeries® and IBM RS/6000® User's Guide.
4. From the System Management Interface Tool (SMIT), configure and start Service Agent Connection Manager (SACM). The SACM is responsible for establishing connectivity to service and support. It enables the gateway server and clients to use a single, secure connection.
   a. Verify that the host name for the SACM is correct. In this example, the SACM resides on the server or logical partition with the modem. Therefore, the SACM is configured to the host name of the server or logical partition with the modem.
   b. Verify the default port 1198. In most cases, the default port is appropriate. You can change the port later, if necessary. This port is necessary for communication between the gateway server and the SACM.
5. Configure and start the Electronic Service Agent gateway server. This is the server or logical partition that acts as the central management server for all of the clients (monitored servers or logical partitions). The Service Agent gateway server contains the central database, and it initiates communication to service and support. The Service Agent gateway communicates to service and support through the SACM.
   a. Verify that the host name is correct. In this example, the SACM and SA gateway server are located on the same server. It is the server or logical partition with the modem. The SA gateway server is the server or logical partition with the modem. Therefore, the SA gateway is configured to the host name of the server or logical partition with the modem.
   b. Verify that the machine type, model, and serial number are correct.
6. Configure and start the Electronic Service Agent clients. This is the monitored server or logical partition for which system information is collected and reported to service and support.
   a. Verify that the host names for the SA clients and for the SA gateway are correct.
   b. Verify that the machine type, model, and serial number are correct.

7. Continue with "Task 12. Use the Service Agent (SA) Basic User Interface."
8. To configure Linux, review the following information.

   In this example, you will configure the following:

   - Electronic Service Agent on the server that has a modem for dial-up connection to service and support

     Note: Alternatively, you can use an Internet or Secure Sockets Layer (SSL) connection instead of the modem to connect to service and support.
   - Electronic Service Agent on the other clients to communicate with the server that has the modem

Figure 10. This diagram shows three servers and their connection through a modem to service and support.

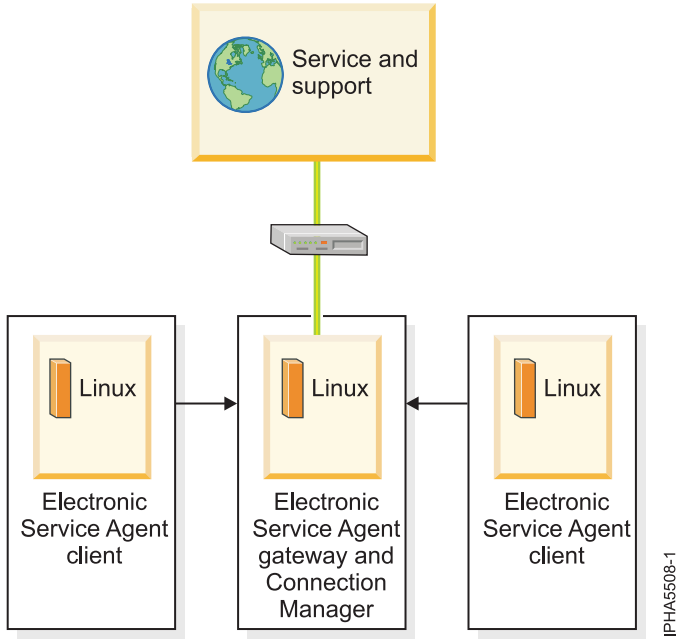

9. From the Linux command line, type the following command to configure and start Service Agent Connection Manager (SACM):

   ```
   startsrc -s sacm
   ```

   The SACM application enables the gateway server and client servers to use a single, secure connection to reach service and support.
10. From the Linux command line, type the following command to configure and start the Electronic Service Agent gateway server:

    ```
    /usr/svcagent/bin/sagatewayconfig
    ```

    The Service Agent gateway server acts as the central management server for all of the clients (monitored servers or logical partitions). It contains the central database and initiates communication to service and support.
11. Continue with "Task 12. Use the Service Agent (SA) Basic User Interface."

Task 12. Use the Service Agent (SA) Basic User Interface

You will need to install Electronic Service Agent on AIX or Linux to access the SA Basic User Interface. For details, refer to the Electronic Service Agent (http://www.ibm.com/support/electronic/navpage?category=5) Web site and search for the appropriate Electronic Service Agent user's guide.

Note: Reference codes based on events detected by the InfiniBand switches (CBxxxxxxx) do not support the Call Home feature. They will not be called home by Electronic Service Agent. If you attempt to use the manual function to initiate Call Home, it will not successfully generate a service call, and you will not receive any indication that the service call was not generated.

While the IBM Network Manager reported serviceable events from an InfiniBand switch will not call home, you will want to consider connectivity to service and support for serviceable events and inventory tools for servers.

1. Familiarize yourself with the SA Basic User Interface. The SA Basic User Interface provides a list of properties and the associated fields that you need to complete to configure Electronic Service Agent.
2. Specify information for the required property fields. Click each property on the left side of the interface, and complete the required fields on the right side of the interface. Required fields are indicated with an exclamation point.

   Depending on how you complete the fields, the interface guides you through the appropriate properties. For example, if you specify that you want to use a modem in the ConnectionManager property fields, the interface automatically displays the Dialer property fields, so that you can complete the information about your modem.

   For this example, where you have multiple servers or logical partitions running AIX or Linux and you use a modem for outbound connectivity, you need to complete specific information for the following properties:

   - **ConnectionManager**: Clear **False** for **Connect to SDR using Dialer** to enable the Dialer. This indicates that you want to use a modem to connect to service and support.
   - **Dialer**: Specify details about your modem and service and support connection parameters.
   - **Machines**: Add two SA client servers.
   - **Enroll**: Register the servers with service and support. This initiates a call to service and support to enroll the servers in service and support's database. To complete the process, service and support sends you a key.
   - **Call log**: Check the status of the call to service and support. You can see whether the call to service and support is successful.

   To learn about advanced features that go beyond the scope of this example, go to the Electronic Service Agent (http://www.ibm.com/support/electronic/navpage?category=5) Web site and search for the appropriate Electronic Service Agent user's guide.

Task 13. Register the ID with Electronic Service Agent for AIX or Linux

1. From the Service Agent Basic User Interface, click **Enroll**.
2. Complete the required fields on the right side of the interface. Required fields are indicated with an exclamation point.
3. Continue with "Task 14. Configure the service processor."

Task 14. Configure the service processor

You might use this type of service connection if your server is not available, because the service processor does not require an operating system to perform its tasks.

To set up your service processor to connect to service and support, you need to attach a modem to the system port on your server. In addition, you need to use the Advanced System Management Interface (ASMI) menus to perform several configuration steps.

1. On the ASMI Welcome pane, specify your user ID and password, and click **Log In**.

   Note: To perform these tasks, you must have Administrator or Service provider authority level.
2. In the navigation area, expand **System Service Aids**.

3. To configure the service processor system port, follow these step:
   a. Select **Serial Port Setup**.
   b. Modify the appropriate fields in the S1 (used with the call-home feature) and S2 (used with the call-in feature) sections.
   c. Click **Save settings** to save the setting changes.
4. To configure the modem, follow these steps:
   a. Select **Modem Configuration**.
   b. Modify the appropriate fields in the S1 and S2 sections.
   c. Click **Save settings**.
5. To configure the call-home and call-in policy, follow these steps:
   a. Select **Call-in/Call-home**.
   b. Enter the desired text into the specified fields.
   c. Click **Save settings** to save changes.
6. To test the call-home policy, follow these steps:
   a. Select **Call-Home Test**.
   b. Click **Initiate call-home test**. A test of the call-home system is performed as specified by the current port and modem selections.
7. Continue with "Task 15. Test the connection to service and support."

Task 15. Test the connection to service and support

1. If you use an HMC to connect to service and support, follow these steps to test the connection for the HMC:
   a. On the HMC, open **Service Applications > Service Focal Point > Service Utilities**.
   b. Select a system.
   c. Select **Selected > Create serviceable event**.
   d. Select **Test automatic problem reporting**.
   e. Select **Request Service**. A message is displayed when the service request is sent.
2. To test the connection for AIX (if you set up Electronic Service Agent on the server or logical partition), follow these steps:
   a. From the System Management Interface Tool (SMIT) on your Electronic Service Agent server, activate the Electronic Service Agent.
   b. Ensure that the Electronic Service Agent Connection Manager is active if it resides on a machine other than the Electronic Service Agent server.
   c. From SMIT, start the Service Agent Advanced User Interface.
   d. To use a modem, configure the Dialer on the Connection Manager screen.

      Note: The default is to connect to service and support using an existing Internet connection.
   e. Open the **Manual Tools** folder.
   f. Select **Connect**.
   g. Monitor the CallLog for the following entry: TEST Connection (Success: 1, Fail: 0).
3. To test the connection for Linux (if you set up Electronic Service Agent on the server or logical partition), follow these steps:
   a. On your Electronic Service Agent server, activate the Electronic Service Agent.
   b. At a Linux command line, type the following:

      ```
      startsrc -g svcagent
      ```
   c. Ensure that the Electronic Service Agent Connection Manager is active if it resides on a machine other than the Electronic Service Agent server.

d. At a Linux command line, type the following:
   ```
   startsrc -s sacm
   ```
e. Start the Service Agent Advanced User Interface.
f. At a Linux command line, type the following:
   ```
   /usr/svcagent/bin/sauiascii
   ```
g. If you want to use a modem, configure the Dialer on the Connection Manager screen.

   Note: The default is to connect to service and support using an existing Internet connection.
h. Open the **Manual Tools** folder.
i. Click **Connection**.
j. Monitor the CallLog for the following entry: TEST Connection (Success: 1, Fail: 0).

Task 16. View the server information that was reported to IBM

IBM provides service and support on the Internet where you can view details of the system you have enabled, and use the data collected by Service Agent. To use the advanced features and receive the full benefits of Electronic Service Agent, you must enter an IBM Registration ID (IBM ID). The first IBM ID entered will have Administrator authority and is able to authorize additional users on the Web site. The second IBM ID is available as a backup for the Administrator.

1. Go to the IBM Electronic Services (http://www.ibm.com/support/electronic) Web site.
2. Click **Sign in** (in the upper right corner).
3. Type the IBM ID and password.
4. Choose the following options from the navigation bar:
   - Select **My systems** to view your server information.
   - Select **Premium Search** to search technical support using your server information to improve the search results.

     Note: In some cases, the Premium Search feature is available only while your server is under warranty or afterward through a service contract.
   - Select **My messages** to view information related to your servers and interests.
5. Enter the requested information.

Getting fixes for a clustered environment

Learn how to get fixes for the clustered server.

When you use a clustered environment, you might need to apply several types of fixes. The fixes include server firmware, HMC (if you are using one), and I/O adapter and device fixes.

For information about fixes you want to apply, refer to the *Fixes and Upgrades*topic in the IBM Systems Hardware Information Center.

Service information and procedures for networks managed by the IBM Network Manager

This section contains a troubleshooting table linking you to reference documentation and problem isolation procedures that are used when servicing GX HCA only and mixed PCI and GX HCA-based InfiniBand cluster networks.

Note: IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.

If you use the IBM Network Manager to service InfiniBand cluster networks, use the following troubleshooting table to link to the appropriate reference documentation or isolation procedures. If you are not using IBM Network Manager, refer to "InfiniBand switch reference information" on page 3.

Task	Procedure Reference
Finding the other side of the link	"Finding the other side of the link" on page 122
Running link diagnostics	"Link diagnostic procedures" on page 165
Removing and replacing an adapter card **Note:** After repair procedures, verify network function according to the "Repair verification" on page 150 procedure.	Refer to "Installing or replacing a GX Host Channel Adapter" on page 64 **Note:** If you are considering deferred maintenance of the adapter, review "Deferring replacement of a failing Host Channel Adapter" on page 66.See the appropriate server repair information in the IBM Systems Hardware Information Center.
Removing and replacing a switch card **Note:** After repair procedures, verify network function according to the "Repair verification" on page 150 procedure.	See the appropriate switch repair manual listed in "InfiniBand switch reference information" on page 3.
Missing LIMs for a 7048-270 or SFS7008P from the Switch Topology View	1. First check the seating of the LIMs and see if their LEDs are properly lit as indicated in the "Interpreting LEDs" on page 122 section. 2. Make sure that the LIMs adhere to the plugging rules in "Plugging Rules for 7048-270 or SFS7008P switch line interface modules and fabric controllers" on page 152
Interpreting various LEDs	"Interpreting LEDs" on page 122
Isolating a problem on a port or a link	"Isolating a problem with a port or link" on page 136
Isolating performance problems	"Isolating performance problems" on page 138
Isolating switch power problems	"Suspected power problem" on page 139
Isolating switch thermal problems	"Suspected thermal problem" on page 140
Verifying that adapters are configured and available	"Verifying adapters are configured and available" on page 140
Checking the subnet manager	"Checking the subnet manager" on page 140
Reordering subnet manager priority	"Reordering Subnet Manager priority" on page 142
Synchronizing the subnet manager time with the HMC time	"Synchronizing Subnet Manager time with HMC time" on page 142

© Copyright IBM Corp. 2004, 2007

Task	Procedure Reference
Updating switch software	“Checking switch software levels” on page 144
Handling time-stamp differences between subnet manager, HMCs, and other devices on the service network	“Understanding timestamp differences” on page 145
Setting GID-prefixes	“Setting GID prefixes” on page 147
Checking GID-prefixes	“Checking GID-prefixes” on page 146
LID Mask Control (LMC) procedures	“Logical identifier mask control procedures” on page 147
Setting LMC	“Setting the location identifier mask control” on page 148
Checking LMC	“Checking the logical identifier mask control” on page 147
Rebooting the entire switch chassis	“Rebooting the entire switch chassis” on page 145
Recovering from an HCA preventing a logical partition from activating	“Recovering from an HCA preventing a logical partition from activating” on page 100
Other service procedures requiring access to the switch management port	“Administrative procedures for InfiniBand switches” on page 157

Reference codes for GX HCA-based InfiniBand networks

Reference code information for InfiniBand (IB) diagnostic procedures.

Information in this section is for diagnostic procedures related to InfiniBand network hardware. Use the following set of symptoms to help determine which actions to take:

Note:

- After repair procedures, verify network function according to the “Repair verification” on page 150 procedure.
- Reference codes based on events detected by the InfiniBand (IB) switches (CBxxxxxxx) do not support the Call Home feature. They will not be called home by Electronic Service Agent. If you attempt to use the manual function to initiate Call Home, it will not successfully generate a service call, and you will not receive any indication that the service call was not generated.

Reference code	Event name	Isolation procedure
CB102800	authenticationFailure	“IBNSAUT” on page 105
CB108800	tsCardDown	“IBNSREM” on page 115
CB109800	tsPowerSupplyDown	“IBNSPOW” on page 113
CB10A000	tsFanDown	“IBNSFAN” on page 110
CB10B800	tsCardRemove	“IBNSREM” on page 115
CB10C100	tsEvent::ibSmSlaveToMaster(4)	“IBNSNLS” on page 113
CB10C1C0	tsEvent::ibSmNodeDeleted(7)	“IBNSNLS” on page 113
CB10C4C0	tsEvent::ibSmDbSyncNotSupported(19)	“IBNSNLS” on page 113

Reference code	Event name	Isolation procedure
CB10C500	tsEvent::ibSmDbSyncNotEnabled(20)	"IBNSDBS" on page 107
CB10C540	tsEvent::ibSmDbSyncNoStandby(21)	"IBNSNLS" on page 113
CB10C580	tsEvent::ibSmDbSyncDbVersionMismatch(22)	"IBNSSMU" on page 119
CB10C5C0	tsEvent::ibSmDbSyncColdSyncTimeout(23)	"IBNSSMR" on page 118
CB10C600	tsEvent::ibSmDbSyncSessionTimeout(24)	"IBNSDBT" on page 108
CB10C641	tsEvent::hardwareError(25)::local I2C error::1	"IBNSDIG" on page 109
CB10C642	tsEvent::hardwareError(25)::remote I2C error::2	"IBNSDGA" on page 108
CB10C643	tsEvent::hardwareError(25)::card seeprom error::3	"IBNSDIG" on page 109
CB10C644	tsEvent::hardwareError(25)::DiskOnChipError::4	"IBNSDIG" on page 109
CB10C645	tsEvent::hardwareError(25)::single bit memory error::5	"IBNSDIG" on page 109
CB10C646	tsEvent::hardwareError(25)::double bit memory error::6	"IBNSDIG" on page 109
CB10C647	tsEvent::hardwareError(25)::real-time clock stopped::7	"IBNSBAT" on page 105
CB10C648	tsEvent::hardwareError(25)::real-time clock sync::8	"IBNSBAT" on page 105
CB10C649	tsEvent::hardwareError(25)::no fan error::9	"IBNPFAN" on page 182
CB10C64A	tsEvent::hardwareError(25)::FPGA error::10	"IBNSDIG" on page 109
CB10C64B	tsEvent::hardwareError(25)::IB switch ASIC error::11	"IBNSDIG" on page 109
CB10C64C	tsEvent::hardwareError(25)::IB switch firmware error::12	"IBNSDIG" on page 109
CB10C64D	tsEvent::hardwareError(25)::voltage/current error::13	"IBNSPOW" on page 113
CB10C64E	tsEvent::hardwareError(25)::cpu ext bus error::14	"IBNSDIG" on page 109
CB10C64F	tsEvent::hardwareError(25)::cpu bus error::15	"IBNSDIG" on page 109
CB10C650	tsEvent::hardwareError(25)::fru error::16	"IBNSDIG" on page 109
CB10C651	tsEvent::hardwareError(25)::local ethernet error::17	"IBNSDIG" on page 109
CB10C652	tsEvent::hardwareError(25)::management ethernet error::18	"IBNSDIG" on page 109
CB10C653	tsEvent::hardwareError(25)::cpu error::19	"IBNSDIG" on page 109

Reference code	Event name	Isolation procedure
CB10C654	An InfiniBand switch card has reported an over temperature condition.	"IBNSTHM" on page 120
CB10C655	tsEvent::hardwareError(25)::na::21	"IBNSSWE" on page 119
CB10C656	An InfiniBand switch card has reported a GUID error.	"IBNSDIG" on page 109
CB10C657	An InfiniBand switch card is not being detected properly.	"IBNSDIG" on page 109
CB10C658	An InfiniBand switch card sensor is not responding properly to a command.	"IBNSDIG" on page 109
CB10C659	An InfiniBand switch card has detected a card fault.	"IBNSDIG" on page 109
CB10C65A	tsEvent::hardwareError(25)::na::26	"IBNSSWE" on page 119
CB10C65B	tsEvent::hardwareError(25)::na::27	"IBNSSWE" on page 119
CB10C65C	tsEvent::hardwareError(25)::na::28	"IBNSSWE" on page 119
CB10C65D	tsEvent::hardwareError(25)::na::29	"IBNSSWE" on page 119
CB10C65E	tsEvent::hardwareError(25)::na::30	"IBNSSWE" on page 119
CB10C65F	tsEvent::hardwareError(25)::na::31	"IBNSSWE" on page 119
CB10C660	tsEvent::hardwareError(25)::na::32	"IBNSSWE" on page 119
CB10C661	tsEvent::hardwareError(25)::na::33	"IBNSSWE" on page 119
CB10C662	tsEvent::hardwareError(25)::na::34	"IBNSSWE" on page 119
CB10C663	tsEvent::hardwareError(25)::na::35	"IBNSSWE" on page 119
CB10C664	tsEvent::hardwareError(25)::na::36	"IBNSSWE" on page 119
CB10C665	tsEvent::hardwareError(25)::na::37	"IBNSSWE" on page 119
CB10C666	tsEvent::hardwareError(25)::na::38	"IBNSSWE" on page 119
CB10C667	tsEvent::hardwareError(25)::na::39	"IBNSSWE" on page 119
CB10C6C1	tsEvent::softwareInitiatedReboot(27); watchdog::1	"IBNSSWP" on page 120
CB10C6C2	tsEvent::softwareInitiatedReboot(27); assertion::2	"IBNSSWP" on page 120
CB10C6C3	tsEvent::softwareInitiatedReboot(27); out-of-memory::3	"IBNSSWP" on page 120
CB10C6C4	tsEvent::softwareInitiatedReboot(27); no working fan::4	"IBNPFAN" on page 182

Reference code	Event name	Isolation procedure
CB10C701	tsEvent::hardwareInitiatedReboot(28); watchdog::1	"IBNSSWP" on page 120
CB100001	tsSensor::(tsSensor::tsDevSensorTemperature=warning)::1	"IBNSTHM" on page 120
CB100002	tsSensor::(tsSensor::tsDevSensorTemperature=normal)::2	"IBNSTHM" on page 120
CB100003	tsSensor::(tsSensor::tsDevSensorTemperature=critical)::3	"IBNSTHM" on page 120
CB10F001	Lost communication with a switch	"IBNSSLC" on page 116
CB10F003	IBM Network Manager cannot register for traps with a switch	"IBNSREG" on page 114
CB10F004	No master Subnet manager found on subnet	"IBNSSMM" on page 118
CB10F005	Timeout on InfiniBand switch explore	"IBNSEXP" on page 110
CB10FF00	VPD problem	"IBNSVPD" on page 121
CB10FF10	POST/Runtime Error on a FRU	"IBNSDIG" on page 109
CB201800	linkDown	"IBNSLNK" on page 111
CB20C230	tsEvent::ibPmConnectionMonitorUtilChange(32)	"IBNSLNK" on page 111
CB20C740	tsEvent::ibPmPortMonitorThresholdErrorChange(29)	"IBNSLNK" on page 111
CB20C780	tsEvent::ibPmPortMonitorThresholdUtilChange(30)	"IBNSLNK" on page 111
CB20C7C0	tsEvent::ibPmConnectionMonitorErrorChange(31)	"IBNSLNK" on page 111
CBFF0000	Network Manager Diagnostics Event - Report this side	"IBNNMD" on page 102
CBFF0001	Network Manager Diagnostics Event - Report other side	"IBNNMD" on page 102
CBFF0002	Network Manager Diagnostics Event - Report cable	"IBNNMD" on page 102
CBFF00FF	Unrecognized Unit Model	"IBNNURM" on page 102
CBFFFF00	Invalid InfiniBand network configuration	"IBNSCFG" on page 106
CBFFFFFE	Network Event : Invalid InfiniBand network configuration	"IBNSCFG" on page 106
CBFFFFFF	Network Event : Invalid InfiniBand network configuration	"IBNSCFG" on page 106

Miscellaneous service procedures for networks managed by the IBM Network Manager

When servicing GX HCA-based InfiniBand cluster networks that are managed by the IBM Network Manager, locate a problem in the troubleshooting table and go to the appropriate reference documentation or problem isolation procedures.

If you are not using IBM Network Manager, refer to "InfiniBand switch reference information" on page 3.

Notes:

1. Before making repairs, review the section, "Repairs when using IBM Network Manager" on page 148 and do the "Repair preparation" on page 148 procedure. After repair procedures, verify network function according to the "Repair verification" on page 150 procedure.
2. For non-7048 machine types, the IBM service representative's responsibilities for repair and isolation actions stop with the InfiniBand switch. Once the IBM service representative has isolated a problem to a switch or a component within a switch chassis, service responsibility is typically transferred to another party.

The following table lists the InfiniBand service procedures.

Task	Procedure Reference
Diagnosing an InfiniBand switch that will not boot	"Diagnosing an InfiniBand switch that will not boot" on page 101
PerformingInfiniBand isolation procedures	"InfiniBand problem isolation procedures" on page 101
Finding the other side of the link	"Finding the other side of the link" on page 122
Interpreting LEDs	"Interpreting LEDs" on page 122
Manipulating FRU identification LEDs	"Manipulating FRU identification LEDs" on page 136
Isolating a problem on a port or a link	"Isolating a problem with a port or link" on page 136
Isolating performance problems	"Isolating performance problems" on page 138
Isolating switch power problems	"Suspected power problem" on page 139
Isolating switch thermal problems	"Suspected thermal problem" on page 140.
Verifying that adapters are configured and available	"Verifying adapters are configured and available" on page 140
Checking the subnet manager	"Checking the subnet manager" on page 140
Reordering the subnet manager priority	"Reordering Subnet Manager priority" on page 142
Synchronizing the subnet manager time with the HMC time	"Synchronizing Subnet Manager time with HMC time" on page 142
Use this procedure when servicing an InfiniBand switch that is not a 7048 machine type.	"Location codes for machine types other than 7048" on page 143
Checking switch software levels	"Checking switch software levels" on page 144
Updating switch software levels	Refer to "Updating switch software" on page 144.
Recovering from a switch being assigned to the incorrect subnet	Refer to "Switch is on the incorrect subnet" on page 144.
Handling time-stamp differences between subnet manager, the HMCs, and other devices on the service network	"Understanding timestamp differences" on page 145

© Copyright IBM Corp. 2004, 2007

Task	Procedure Reference
Recovering from logical HCA configuration problems	Refer to “Recovering from logical HCA configuration problems” on page 145.
Rebooting the entire switch chassis	“Rebooting the entire switch chassis” on page 145
Adjusting firewall parameters for SNMP traps	“Adjusting Firewall Parameters for SNMP Traps” on page 146
Setting GID-Prefixes	“Setting GID prefixes” on page 147
Checking GID-Prefixes	“Checking GID-prefixes” on page 146
LID Mask Control (LMC) Procedures	“Logical identifier mask control procedures” on page 147
Setting LMC	“Setting the location identifier mask control” on page 148
Checking LMC	“Checking the logical identifier mask control” on page 147
Running link diagnostics	“Link diagnostic procedures” on page 165
Removing and replacing an adapter	Refer to “Installing or replacing a GX Host Channel Adapter” on page 64 **Note:** If you are considering deferred maintenance of the adapter, review “Deferring replacement of a failing Host Channel Adapter” on page 66.Refer to the appropriate server repair manual listed in “InfiniBand switch reference information” on page 3.
Remove and replace a switch card	Refer to the appropriate InfiniBand switch repair documentation listed in “InfiniBand switch reference information” on page 3.
Examine and collect the ts_log file from a switch	Refer to “Reviewing the ts_log file” on page 159.
Repairs when using the IBM Network Manager	“Repairs when using IBM Network Manager” on page 148
Check the Subnet Manager	“Checking the subnet manager” on page 140
Setting IP addressing in switches	“Setting IP addressing in switches” on page 159
Disappearing IBM Network Manager Windows	“Disappearing IBM Network Manager windows” on page 152
Verifying static-12x or 4x configuration for a port	“Verifying Static12x or 4x configuration to a port” on page 153
Recovering from an HCA preventing a logical partition from activating	“Recovering from an HCA preventing a logical partition from activating”
Do other service procedures that require access to the switch management port	“Administrative procedures for InfiniBand switches” on page 157

Recovering from an HCA preventing a logical partition from activating

Use this procedure to recover a logical partition when a failed HCA is preventing the partition from activating.

During IPL a logical partition can be prevented from activating because an HCA has failed. Do the following to unassign HCAs from partition profiles.

1. Go to the **Server and Partition** window.
2. Click the **Server Management** partition.
3. Expand the server in which the HCA is installed.
4. Expand the partitions under the server.

5. Do the following procedure for each partition profile that uses the HCA. If you do not know which partitions use the HCA, you must perform the following procedure for each partition profile:
   a. Select each partition profile that uses the HCA.
   b. From the menu, click **Selected → Properties**.
   c. In the Properties dialog, click the **HCA** tab.
   d. Using its physical location, find the HCA of interest.
   e. Highlight the HCA of interest and then click the Clear button. The HCA's GUID Index, GUID, and Capability fields change to `Unassigned`. Then click the **Ok** button.

      Note: A failing HCA can be unassigned from the logical partition profile while the logical partition is active, hung, or inactive. If the logical partition is currently active, the logical partition needs to be shutdown and then activated for this update to take effect. You do not have to do reactivate the logical partition if you are deferring maintenance on an HCA. By changing the defective HCA to "Unassigned" in the partition profile, you are ensuring that the next activation is not prevented by a failing HCA.
6. **This procedure ends here.**

Diagnosing an InfiniBand switch that will not boot

Use this procedure to diagnose a problem with a switch that will not boot.

Typically, if a switch will not boot, the power status for the switch is one of the following:

- Power transition
- Power IPL transition
- Power dump

This power status is found in the Switch Topology View window on the IBM Network Manager interface. See "Switch topology window status" on page 200.

If a switch does not boot, do the following:

1. Check the Service Focal Point on the HMC that is running the IBM Network Manager.
2. Check the Chassis/System-Wide LEDs. If the Chassis/System-Wide LED indicates a problem, perform the appropriate service procedure as outlined in "Switch chassis and system-wide LEDs" on page 124.
3. Inspect all other LEDs on the chassis, and perform actions as prescribed in "Interpreting LEDs" on page 122.
4. Verify that you have an appropriate level of switch-management software on this switch; see "Checking switch software levels" on page 144. Correct the level of software if necessary.
5. If the previous actions have not resolved the problem, contact your next level of support.
6. **This procedure ends here.**

InfiniBand problem isolation procedures

Problem isolation procedures used for InfiniBand switch errors.

For InfiniBand network errors, use the problem isolation procedures in this section when you are directed to do so.

Note: After repair procedures, verify network function according to the "Repair verification" on page 150 procedure.

The following list identifies the InfiniBand problem isolation procedure in alphabetic order.

"IBNNMD"	"IBNSCRD" on page 107	"IBNSDGA" on page 108	"IBNSNLS" on page 113	"IBNSSMU" on page 119
"IBNNURM"	"IBNSCTL" on page 185	"IBNSDIG" on page 109	"IBNSPOW" on page 113	"IBNSSWE" on page 119
"IBNSAUT" on page 105	"IBNSDBS" on page 107	"IBNSFAN" on page 110	"IBNSREM" on page 115	"IBNSSWP" on page 120
"IBNSBAT" on page 105	"IBNSDBT" on page 108	"IBNSFRU" on page 111	"IBNSSLC" on page 116	"IBNSTHM" on page 120
		"IBNSLNK" on page 111	"IBNSSMR" on page 118	"IBNSVPD" on page 121

Note: Reference codes based on events detected by the InfiniBand switches (CBxxxxxxx) do not support the Call Home feature. They will not be called home by Electronic Service Agent. If you attempt to use the manual function to initiate Call Home, it will not successfully generate a service call, and you will not receive any indication that the service call was not generated.

IBNNMD

Use this procedure if the user was running diagnostics as an isolation method for a link that is suspected as being bad.

You are here because you have noted this isolation procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx.

Note: After repair procedures, verify network function according to the "Repair verification" on page 150 procedure.

This indicates that a user has selected to create a serviceable event based on the results of link diagnostics. This should have been done if the user was running diagnostics as an isolation method for a link that is suspected as being bad, but no serviceable event existed to isolate to the tested link. This is an effective method for keeping track of service in a cluster. Perform the following procedure:

Note: IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.

1. If nothing has been replaced yet, replace the part indicated in the FRU list. Otherwise, proceed to the next step.
2. Because this was generated for tracking purposes, carefully record your actions in the serviceable event. Note the replaced FRU and make any comments that you deem useful for tracking purposes.
3. **This procedure ends here.**

IBNNURM

Use this isolation procedure when an error has been reported against a switch that has an unrecognized model.

You are here because you have noted this isolation procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx.

Note:

- IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service is responsible.
- After repair procedures, verify network function according to the "Repair verification" on page 150 procedure.

This isolation procedure is indicated because an error was reported for a switch that, according to the IBM Network Manager, has an unrecognized model. The IBM Network Manager cannot be used to isolate the FRU. It is possible that you will need to transfer this service activity to another service organization.

IBNNURM procedure overview

This procedure helps you determine if a switch is really an unrecognized switch model, or if there is a failure in the switch that makes it unrecognizable. The basic steps are:

1. Record information and interpret the unit location code.
2. Check the label on the switch to see if the model matches what IBM Network Manager is displaying.
3. If the label indicates that this is not an IBM 7048-120 or 7048-270 InfiniBand switch, work with the customer to identify the switch manufacturer and service provider.
4. Contact the appropriate service provider.

Determine the correct service provider for this switch

To determine the correct service provider for this switch, perform the following procedure:

1. Record the following information about the event:
   - Reference code
   - FRU information
   - FRU location code for the switch in the FRU list. This should be the first item with the class of FRU.

     Important: The unit location portion of the location code is displayed as follows:
     U*unit_location*-P1-*remainder_of_the_location_code*
2. Use the following table to cross-reference the unit location format to the machine type for the switch that is in the FRU list.

Unit location format	Machine type of switch	Instructions
U7048.120.*serial_number*	7048 known model	Proceed to collect data and call your next level of support.
U7048.270.*serial_number*	7048 known model	Proceed to collect data and call your next level of support.
USFS7000Psssssss.*serial_number*	SFS7000P known model	Proceed to collect data and call your next level of support.
USFS7008Psssssss.*serial_number*	SFS7008P known model	Proceed to collect data and call your next level of support.
U7048.###.*serial_number*	7048 unknown model	Proceed to the correct procedure.
U7048.###.*sssssss*	7048 unknown model	Proceed to collect data and call your next level of support.
U*tttt*.###.*sssssss*	Unknown	Proceed to collect data and call your next level of support.

Unit location format	Machine type of switch	Instructions
U*.* (the * can be any character)	Unknown	Proceed to collect data and call your next level of support.
UTS120*sssssss*.US*serial_number*	IBM System x™, Topspin, or Cisco	Go to the instructions for IBM System x, Topspin, or Cisco switches.
UTS270*sssssss*.US*serial_number*	IBM System x, Topspin, or Cisco	Go to instructions for IBM System x, Topspin, or Cisco switches.

3. If the machine type of the switch is a known 7048 machine, use the following procedure. Otherwise, go to the next step.

   Note: In this case, there is likely a software issue that is causing IBM Network Manager to incorrectly determine that it cannot recognize the model number of this switch.

   a. Log in to the HMC and:
      1) Run: `/opt/hsc/bin/ibnm.snap`

         The output of ibnm.snap will be in `/var/hsc/log/[HMCname].[timestamp].snap.tar.gz`
      2) Save the newest extended error data found in: `/var/hsc/log/[reference code]/*.snap.gz`
   b. Call your next level of support.
   c. **This procedure ends here**.
4. If the machine type and model of the switch is not known but the serial number is known, use the following procedure. Otherwise, proceed to the next step.
   a. Match the serial number in the location code with the chassis serial number on the back of a switch.
   b. If the matching switch is a 7048-120 or SFS7000P, perform the following procedure:
      1) Check the comments for the serviceable event. If this has previously happened to this switch, replace the switch chassis. Otherwise, proceed to the next step.
      2) Reboot the switch chassis using the procedure found in “Rebooting the entire switch chassis” on page 145.
      3) Go to the Switch Topology view in the IBM Network Manager, and find this switch using the serial number.
      4) If the location code for the switch still does not have a valid model number in the unit location (U7048.120 (or 700).[*serialNumber*], replace the switch chassis.
      5) If the location code for the switch has a valid model number in the unit location, update the comments for the serviceable event indicating that a reboot fixed issue, and then close the event.
      6) **This procedure ends here**.
   c. If the matching switch is a 7048-270 or SFS7008P, perform the following procedure:
      1) Check the comments for the serviceable event. If this has previously happened to this switch, replace the switch chassisID module in slot 17. If the chassisID module has already been replaced, replace the master Fabric Controller Module. To find the master Fabric Controller Module, see “Fabric controller LEDs” on page 130. Otherwise, proceed to the next step.
      2) Reboot the switch chassis, using the procedure found in “Rebooting the entire switch chassis” on page 145.
      3) Go to the Switch Topology view in the IBM Network Manager GUI, and find this switch using the serial number.
      4) If the location code for the switch still does not have a valid model number in the unit location (U7048.270 (or 708).[serialNumber], replace the switch chassisID module in slot 17. If the chassisID module has already been replaced, replace the master Fabric Controller Module. To find the master Fabric Controller Module, see “Fabric controller LEDs” on page 130.

5) If the location code for the switch has a valid model number in the unit location, update the comments for the serviceable event indicating "reboot fixed issue," and close the event.
6) **This procedure ends here.**

d. If there is no matching switch found, there is a software problem. Perform the following procedure:
1) Log in to the HMC and:
a) Run: `/opt/hsc/bin/ibnm.snap`
The output of ibnm.snap will be in `/var/hsc/log/[HMCname].[timestamp].snap.tar.gz`
b) Save the newest extended error data found in: `/var/hsc/log/[reference code]/*.snap.gz`
2) Call your next level of support.
3) **This procedure ends here.**

5. If the machine type of the switch is 7048 unknown model and the serial number is not known, use the following procedure. Otherwise, proceed to the next step.
a. Log in to the HMC and:
b. Run: `/opt/hsc/bin/ibnm.snap`
The output of ibnm.snap will be in /var/hsc/log/[HMCname].[timestamp].snap.tar.gz
c. Save the newest extended error data found in: `/var/hsc/log/[reference code]/*.snap.gz`
d. Call your next level of support.
e. **This procedure ends here**.
6. If the Machine Type of the switch is xSeries or Topspin, use the following procedure.
a. Confer with the customer to determine if the original purchase of the switch was from IBM System x, Topspin, or Cisco.
b. If it is an IBM System x switch, direct the customer to the IBM System x support line for service.
c. If it is a Topspin or Cisco switch, direct the customer to the Topspin or Cisco support line for service.
d. **This procedure ends here**.

IBNSAUT

Someone has been trying to access the subnet manager without proper authorization.

You should be here because you have noted this isolation procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx.

Note: IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.

Note: After repair procedures, verify network function according to the "Repair verification" on page 150 procedure.

Someone has been trying to access the subnet manager without proper authorization.

1. If a service action is the likely cause for this, close the event. Otherwise, proceed to the next step.
2. Perform a security audit to determine where this security breach is occurring. Some potential access points are the switch chassis serial port or the service network's Ethernet network.
3. **This ends the procedure**.

IBNSBAT

Isolate a battery failure on a switch.

You are here because you have noted this isolation procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx.

Note:

- IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service is responsible.
- For instructions on FRU replacement, refer to the appropriate repair manual using the references in "InfiniBand switch reference information" on page 3.
- After repair procedures, verify network function according to the "Repair verification" on page 150 procedure.

1. Run card diagnostics on the switch card listed in the FRU list that contains the battery. For instructions on running switch card diagnostics, refer to "InfiniBand switch card diagnostics" on page 173.
2. If diagnostics do not indicate a failure, try rebooting the card in the FRU list that contains the battery. For instructions on rebooting a switch card, refer to "Rebooting switch cards" on page 160.
3. If the problem recurs, replace the switch card in the FRU list.
4. **This ends the procedure.**

IBNSCFG

IBM Network Manager has detected an invalid configuration.

One of the following possible problems is detected with the configuration.

1. More than one switch in a subnet. If there is a single switch in the FRU list, the configuration issue has to do with having too many switches in a subnet.
2. More than four subnets in a cluster. If there are multiple switches in the FRU list, the configuration issue has to do with having too many subnets in the network.

You are here because you have noted this isolation procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx.

Notes:

1. IBM service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.
2. After repair procedures, verify network function according to "Repair verification" on page 150.

Perform the following procedure:

1. Record the serviceable event reference code, serviceable event text, and FRU list.
2. If you have been given permission to configure your InfiniBand network with beyond the default configurations, update the configuration file, **/opt/hsc/data/ibnm/valid_config** on the HMC running IBM Network Manager. This should be done on any HMC that you may use to run IBM Network Manager should the main one fail. Change the value(s) next to the appropriate keyword(s). It is advisable to only give the maximum numbers you will require for your particular configuration so that IBM Network Manager can check for inconsistencies in your configuration. The default values are given below:
   - max_switches_per_subnet 1
   - max_subnets 4

3. If you believe that have a valid configuration, double check cabling and seating of cables between switches. If there are enough cabling issues, it might appear to IBM Network Manager as if switches are not connected when you intended to have them connected. This should only occur if the max_switches_per_subnet value has been changed from the default of greater than one.
4. Call your next level of support and supply them with the above recorded information.
5. **This ends the procedure.**

IBNSCRD

Use this procedure if your serviceable event begins with CBxxxxxx.

You are here because you have noted this isolation procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx.

Note:

- IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.
- For instructions on FRU replacement, refer to the appropriate repair manual using the references in "InfiniBand switch reference information" on page 3.
- After repair procedures, verify network function according to the "Repair verification" on page 150 procedure.

1. Determine if service is required:
   a. Check the duplicate count in the Serviceable Event Details window. If it is greater than 0, service is required.
   b. If the duplicate count is 0, check the comments for the serviceable event.
   c. If the comments for the serviceable event indicate that card diagnostics have been run previously, service is required.
   d. If the comments for the serviceable event do *not* indicate that card diagnostics have been run previously, you may run diagnostics on the card prior to replacing it. For information on running diagnostics, see "InfiniBand switch card diagnostics" on page 173.

      Note: Running card diagnostics temporarily causes the card to stop functioning correctly. This might result in false serviceable events being reported by other devices interfacing with the card. After running diagnostics, wait 5 to 10 minutes for the diagnostics to flow through the service subsystem and then check for false serviceable events that might be caused by this action (especially Link Down events). Close out the false serviceable events as appropriate. Also, note in the comments for the false serviceable events that they were caused by running card diagnostics.

   e. If diagnostics do not fail, you might decide not to replace the card, but note in the serviceable event's comments that the card diagnostics have been run.
2. If service is required, replace the switch card listed in the FRU list for the serviceable event.
3. **This procedure ends here.**

IBNSDBS

A software-initiated reboot has occurred.

You should be here because you have noted this isolation procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx.

Note:

- IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.
- After repair procedures, verify network function according to the "Repair verification" on page 150 procedure.

Do the following:

1. Enable database synchronization on the Infiniband Subnet Manager indicated by the switch card in the FRU list. Refer to "Setting database synchronization" on page 161.
2. **This ends the procedure.**

IBNSDBT

A database synchronization timeout has occurred between the master subnet manager and the highest standby.

You are here because you have noted this isolation procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx.

Note:

- IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.
- After repair procedures, verify network function according to the "Repair verification" on page 150 procedure.

Do the following:

1. Check the master subnet manager IBSM switch and highest standby subnet manager IBSM switch for status and serviceable events related to these devices. For link status interpretation, refer to "Status procedures for the IBM Network Manager" on page 189.
   - If a serviceable event is open against the master subnet manager IBSM switch, perform the prescribed service actions for that serviceable event. If the serviceable event requests you to determine whether it is appropriate to perform service, this request indicates that service is appropriate.
   - If there is no serviceable event open against the master subnet manager IBSM switch or the highest standby subnet manager IBSM switch, check the status of the switch in the IBM Network Manager, and perform any service actions appropriate for the status that you find.
2. Check the link status between the master subnet manager IBSM switch and the highest-standby subnet manager IBSM switch. For link status interpretation, refer to "Status procedures for the IBM Network Manager" on page 189.
   - If there is no active link between the master and highest-standby subnet manager IBSM switch, perform link diagnostics on the link that should be active, and replace FRUs as appropriate.
   - If there is an active link between the master and highest standby subnet manager IBSM switch, run the procedure in "Link diagnostic procedures" on page 165. If no FRU is found, call your next level of service.
3. **This ends the procedure.**

IBNSDGA

An event has occurred that requires card diagnostics to be run on all cards to isolate the FRU.

You are here because you have noted this isolation procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx.

Note:

- IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.
- For instructions on FRU replacement, refer to the appropriate repair manual using the references in "InfiniBand switch reference information" on page 3.
- After repair procedures, verify network function according to the "Repair verification" on page 150 procedure.

1. Run card diagnostics on all cards contained within the switch chassis that are indicated by the first FRU in the FRU list.
2. If diagnostics do not indicate a failure, replace the card that contains the switch controller. This card is the second FRU in the FRU list.
3. **This ends the procedure.**

IBNSDIG

Use this procedure

You are here because you have noted this isolation procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx.

Note:

- IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.
- For instructions on FRU replacement, refer to the appropriate service documentation.
- After repair procedures, verify network function according to the "Repair verification" on page 150 procedure.

Use the following procedure to determine if service is required, and if service is required, isolate to the correct FRU.

1. Record the reference code for the serviceable event.
2. Record the first and last reported times of the serviceable event in the Manage Service Events – Serviceable Event Overview window.
3. Record the duplicate count in the Serviceable event detailed attributes panel in the Manage Serviceable Events – Serviceable Event Details window.
4. Use the following table to determine how to handle the possible reference code for this procedure:

Reference Code	Procedure
CB10C64A, CB10C64B, CB10C64C, CB10C646	1. If the duplicate count is greater than 0, go to 5 on page 110. 2. If the duplicate count is 0, close the serviceable event.
CB10C645	1. If the duplicate count is greater than 15, go to 5 on page 110. 2. If the duplicate count is 15 or fewer, close the serviceable event.

Reference Code	Procedure
All other CBxxxxxx codes	Go to 5.

5. Isolate to the correct FRU.

   Note: For instructions on FRU replacement, refer to the appropriate repair manual using the references in "InfiniBand switch reference information" on page 3.

   a. Run card diagnostics on the switch card in the FRU list. For instructions on running card diagnostics, see "InfiniBand switch card diagnostics" on page 173.
   b. If diagnostics do not indicate a failure, reboot the card and see if the failure recurs. For instructions on rebooting a switch card, refer to "Rebooting switch cards" on page 160.
   c. If the failure recurs, replace the card.
   d. **This ends the procedure.**

IBNSEXP

The maximum time for an explore has been exceeded. The maximum time for an explore is set to 30 minutes. This time limit accounts for potentially busy switches and a busy Ethernet service network.

You are here because you have noted this isolation procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx.

Notes:

1. IBM service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.
2. After repair procedures, verify network function according to "Repair verification" on page 150.

Perform the following procedure:

1. Record the serviceable event reference code, serviceable event text, and FRU list.
2. Check Service Focal Point to see if there are other serviceable events being reported against the switch in the FRU list. If there are, fix those events and close this one.
3. Reboot the switch in the FRU list using the procedure in "Rebooting the entire switch chassis" on page 145.
4. Call your next level of support and supply them with the above recorded information.
5. **This ends the procedure.**

IBNSFAN

A fan event has occurred on a switch.

You are here because you have noted this isolation procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx.

Note:

- IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.
- For instructions on FRU replacement, refer to the appropriate repair manual using the references in "InfiniBand switch reference information" on page 3.

- After repair procedures, verify network function according to the “Repair verification” on page 150 procedure.

Do the following:

1. Check the status of the switch chassis that is cooled by the fan in the serviceable-event FRU list.
2. If the switch chassis is still operational, you can defer maintenance until a later time. Otherwise, proceed to the next step.
3. Replace the fan FRU listed in the FRU list for the serviceable event. If more than one fan FRU is listed in the FRU list, , follow these steps to isolate which fan is failing:
   a. Go to the HMC that is running IBM Network Manager.
   b. From the IBM Network Manager Overview window, choose View Management Properties.
   c. Select the Switch tab.
   d. Highlight the switch called out in the serviceable event.
   e. Select Environmentals.and look for a status for a fan that is other than Normal or Up. For more details on Environmental Status see “Switch Environmental Status” on page 205.
      - If the switch is a 7048-120 or SFS7000P, replace the card located in U[*unit location*]-P1-E*y*; where *y* is 1 for IDs 1 and 2, *y* is 2 for IDs 3 and 4, and *unit location* is U*MachineType*.Model.*SerialNumber* or U*ProductName*.*SerialNumber*; where *ProductName* is SFS7000P* or SFS7008P* .
      - If the switch is a 7048-270 or SFS7008P, replace the card located in U[unit location]-P1-A*y*; where *y* is 1 for IDs 1 and 2, *y* is 2 for IDs 3 and 4, and *unit location* is U*MachineType*.Model.*SerialNumber* or U*ProductName*.*SerialNumber*; where *ProductName* is SFS7000P* or SFS7008P* .

This ends the procedure.

IBNSFRU

Isolate to a FRU for serviceable events that begin with CBxxxxxx.

You are here because you have noted this isolation procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx.

Note:

- IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.
- For instructions on FRU replacement, refer to the appropriate repair manual using the references in “InfiniBand switch reference information” on page 3.
- After repair procedures, verify network function according to the “Repair verification” on page 150 procedure.

Do the following:

1. Replace the FRUs one at a time, in the order in which they are presented in the FRU list, until the event no longer occurs.
2. **This ends the procedure.**

IBNSLNK

Isolate to a FRU for InfiniBand network link problems.

You are here because you noted this problem isolation procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx. When IBNSLNK is listed, it indicates a link failure. Link failures can be caused by a user action, such as disconnecting a cable. If a cable was recently unplugged, reconnect the cable. This is an example of a problem that could have been caused by a user action. If the serviceable event was caused by a user action, correct the cause and then close the serviceable event.

Note:

- IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.
- After repair procedures, verify network function according to the "Repair verification" on page 150 procedure.

If the link failure cannot be attributed to a user action, the service representative must do the following:

1. Record the first and last reported times of the serviceable event in the Manage Service Events – Serviceable Event Overview window.
2. Record the duplicate count in the Serviceable event detailed attributes panel in the Manage Serviceable Events – Serviceable Event Details window.
   a. If the duplicate count is greater than 0, continue to step 3.
   b. If the duplicate count is 0 and there are no perceived performance problems, close the event and wait for it to recur. If there is a perceived performance problem, proceed to step 3.
3. On the HMC, check the serviceable event and the IBM Network Manager.
   a. Record the part numbers and location codes for the cable entries in the FRU list. They are identified in the part number column for the symbolic FRU identifiers IBNCCAB and CBLCONT. These identifiers represent each end of the cable. If only one end of the cable is known, it is likely that the IBM Network Manager does not recognize the device on the other end of the cable.
   b. If there is both an IBNCCAB and a CBLCONT symbolic FRU identifier, then proceed to Isolate to the correct FRU. Otherwise, continue to the next step.
   c. Using the location code for one end of the cable that you recorded above, determine what device is on the other end of the cable. You can do this by:
      - Referring to cable planning information
      - Looking for a label on the cable that indicates what is attached to the other side
      - Following the cable to the other device
   d. If any one of the following is true, then close the event. Then, check the Switch Topology View for good status using the status procedures in "Switch topology window status" on page 200. Otherwise, continue to step Isolate to the correct FRU.
      - The server that contains the other device was recently powered off or rebooted.
      - An error was reported against the other device that might have caused the link to report a problem.
      - A checkstop was reported against the other server that would have caused the link to report an error on reset.
4. Isolate to the correct FRU using the following procedure:

   Note:

   - If there are multiple switch ports in a 7048-270 or SFS7008P that have the same serviceable event against them, you must first attempt to determine if there is a common failure behind them. Use the procedure found in "Determining faulty fabric controller cards versus faulty LIM cards" on page 153. However, first ensure that there is no common event that may have caused this, like rebooting of many servers in the cluster which are connected to this switch card.

- For instructions on FRU replacement, refer to the appropriate repair manual using the references in "InfiniBand switch reference information" on page 3.

a. Check both ends of the cable for obvious loose connections or bent pins, and repair as necessary.

   Note: If this is a 4x link configured to run at 12x with an octopus cable, you may be replacing or reseating a cable with three 4x connectors on one end and one 12x connector on the other end. In such a case, carefully record, which 4x connector goes to which 4x switch port. If you do not do this, you will need to determine the configuration again using the procedure in "Planning for InfiniBand networks" on page 5. For more on behavior of 12x connections see "Planning for InfiniBand networks" on page 5.

b. Run diagnostics on the link indicated by the first FRU in the FRU list. Refer to "Link diagnostic procedures" on page 165.

c. Check both ends of the cable for obvious loose connections or bent pins, and repair as necessary.

d. Replace the cable.

e. Replace the remaining devices in the FRU list, in the order in which they are listed.

   Notes:

   1) Before replacing an HCA, review "Installing or replacing a GX Host Channel Adapter" on page 64. If you are considering deferred maintenance of the adapter, review "Deferring replacement of a failing Host Channel Adapter" on page 66.
   2) If you replaced a switch card or chassis that had octopus cables connected to it, you need to configure all of the ports on the new switch card(s) that are acting as a group in a 12x link using an octopus cable. See the procedure in "Configuring Static-12x groups" on page 20. For more on behavior of 12x connections see "Planning octopus cables in static 12x cabling" on page 19.

5. **This ends the procedure.**

IBNSNLS

The subnet manager in this switch has become the master subnet manager for this InfiniBand subnet.

You should be here because you have noted this isolation procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx.

Note:

- IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service is responsible.
- For instructions on FRU replacement, refer to the appropriate service documentation.
- After repair procedures, verify network function according to the "Repair verification" on page 150 procedure.

1. If this was not a user-initiated change, record the reference code, serviceable event text, and FRU list for this serviceable event.
2. Call your next level of support for further guidance.
3. **This ends the procedure.**

IBNSPOW

A power event has occurred on a switch.

You are here because you have noted this isolation procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx.

Note:

- IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service is responsible.
- After repair procedures, verify network function according to the "Repair verification" on page 150 procedure.

Use this procedure to determine if service is required, and if service is required, isolate to the correct FRU.

1. Check that all power cables are correctly seated. If a cable seating problem was detected, reseat the cable. If no cable seating problems were detected, go to the next step.
2. Record the reference code for the serviceable event.
3. Record the first and last reported times of the serviceable event in the Manage Service Events – Serviceable Event Overview window.
4. Record the duplicate count in the Serviceable event detailed attributes panel in the Manage Serviceable Events – Serviceable Event Details window.
5. If the reference code is CB109800, go to Isolate to the correct FRU..
6. If the reference code is CB10C64D, perform the following procedure:
   a. If the duplicate count is greater than or equal to 11, go to stepIsolate to the correct FRU.
   b. If the duplicate count is less than 11, close the event and wait for it to recur.
7. If the reference code is neither of the above reference codes, call your next level of support.
8. Isolate to the correct FRU:

   Note: For instructions on FRU replacement, refer to the appropriate repair manual using the references in "InfiniBand switch reference information" on page 3.

   a. Check the status of the switch chassis that is powered by the power supply in the FRU list.
   b. If the switch chassis is still operational, you can defer maintenance until a later time. Otherwise, proceed to the next step.
   c. Replace the power FRU listed in the FRU list for the serviceable event. **This ends the procedure.**

IBNSREG

IBM Network Manager cannot register for traps with a switch.

You are here because you have noted this isolation procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx.

Notes:

1. IBM service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.
2. After repair procedures, verify network function according to "Repair verification" on page 150.

1. Record the serviceable event reference code, serviceable event text, and FRU list.
2. Verify that the switch is powered on by going to the switch chassis indicated in the FRU list and inspecting the LEDs. Refer to the "Interpreting LEDs" on page 122.

3. Verify the ethernet connections from the switch indicated in the FRU list to the switch/router on the service ethernet network, as well as the connectivity of the HMC to the service ethernet network. If you can see any other devices on the HMC, it is not likely that the connection from the HMC to the service ethernet network is faulty.
4. Verify that there are no InfiniBand switch to switch connections in your configuration. This is not supported at this time, and can result in unpredictable behavior.
5. Verify that the SNMP port is open on the HMC:
   a. Go to the LAN Adapters panel in the HMC graphical user interface:**HMC management → HMC Configuration → customize network settings → lan adapters**
   b. Select the private interface (usually eth0) and click **Details**.
   c. Click the **Firewall** tab.
   d. In the bottom selection box, you should seen an entry for SNMP traps 162:tcp 162:udp
   e. If you do not see an SNMP traps entry, perform the procedure in "Adjusting Firewall Parameters for SNMP Traps" on page 146. Otherwise, go to the next step.
6. Verify that no other software on the HMC running IBM NM is registered with the switch(es) for traps and listening to IP port 162. For example, the Cisco or Topspin Element Manager.
   a. Check that there isn't anything else listening on port 162:
      1) Open an **xterm** or command line session on the HMC running IBM Network Manager.
      2) Run: **netstat -nlp | grep 162**
      3) If anything other than ibnmd* is listening to port 162, this software must be removed from the HMC. Continue with the next step.
   b. If there is software running on the HMC that can also be registered for switch SNMP traffic, disable and remove that software.
7. Call the next level of software support and supply them with the above recorded information.
8. **This ends the procedure.**

IBNSREM

This procedure can determine if a card has been purposely removed from the switch.

You are here because you have noted this isolation procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx.

Note:

- IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service is responsible.
- For instructions on FRU replacement, refer to the appropriate repair manual using the references in "InfiniBand switch reference information" on page 3.
- If a chassis ID is changed, for example when a VPD module or the switch chassis is replaced, do the following:
  1. Record the original switch chassis serial number.
  2. After replacing the chassis, go to the IBM Network Manager and open the Management Properties window. Verify that the switch has a valid location code. A valid location code for a 7048-120 begins with U7048.120.[*serial number*]. A valid location code for a 7048-270 begins with U7048.270.[*serial number*].
  3. If you have replaced a 7048-120 chassis, remember to use the RID tag procedures to document the original chassis serial number, which is required for service entitlement.

To determine if a card was purposely removed from the switch, use the following procedure.

1. Using the location code of the card in the FRU list, determine if this card was purposely removed.
2. If the card was purposely removed, close out the serviceable event. Otherwise, proceed to the next step.
3. If the card was not purposely removed, reseat the card and see if the event recurs.
4. If the event recurs, replace the card in the FRU list.
5. **This ends the procedure.**

IBNSSLC

Use this procedure when the IBM Network Manager loses communication with a switch.

IBM Network Manager has lost communication with a switch. This can indicate one of the following conditions:

- The switch was powered off or is under repair.
- There is a service subsystem failure.
- There is a failure with the service subsystem function on the switch.

  The service subsystem function in a 7048-120 and SFS7000P is entirely on the switch board. The service subsystem function in a 7048-270 and SFS7008P consists of the core fabric controller modules (slots 11 and 12), and the management I/O modules (slots 15 and 16).
- There is a power problem that caused the switch to go down. The switch may have powered itself off due to both power supplies failing.
- There is a fan/thermal problem that caused the switch to go down. The switch may have powered itself off due to both fans failing.

Note:

- IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service is responsible.
-
  1. If a switch is unresponsive, it is possible that serviceable events may be lost during the time it is unresponsive. Be sure to check the port status of all ports on the switch after it has been fixed.
  2. After repair procedures, verify network function according to the “Repair verification” on page 150 procedure.

The switch chassis is indicated in the location code provided in the FRU list. Perform the following procedure:

1. Determine if the switch is powered off by checking the power status in the IBM Network Manager graphical user interface’s Switch Topology View, or by looking at the chassis power LEDs on the switch itself.
   a. If the switch is powered on, continue on to the next step.
   b. If the switch is not powered on and this is as expected, close the serviceable event and make a comment stating that the “switch was powered off, as expected.” **This procedure ends here.**
   c. If the switch is not powered on, and this is not as expected, power the switch on.
   d. If the switch will not power on, perform the procedure indicated in “Suspected power problem” on page 139. You do not need to check SFP again.
   e. If the power checks out, perform the procedure indicated in “Suspected thermal problem” on page 140. You do not need to check SFP again.
   f. If you do not find a power problem or a thermal problem, call your next level of support.
2. Inspect all Ethernet connections and ports in the service network, to verify that they are operational.

a. If you find a problem with the service network, take the appropriate action to fix it. **This procedure ends here**.
b. If you do not find a problem with the service network, proceed to the next step.

3. If the switch is a 7048-270 or an SFS7008P, perform the following procedure. Otherwise, proceed to the next step.
   a. Gather management I/O module information:
      1) Note if either slot 15 or 16 has a management I/O module installed.
      2) Note the state of the LEDs on each management I/O module . See "Switch management I/O module LEDs" on page 133and perform any actions prescribed.
   b. Gather Fabric Controller core information:
      1) Note if either slot 11 or 12 has fabric controllers installed
      2) Note the state of the LEDs on either core fabric controller (slots 11 and 12. See "Fabric controller LEDs" on page 130and perform any action prescribed.
   c. If no actions have been taken thus far, use the following table to determine which action should be taken based on the Master LED indicators on the Management I/O modules and the core Fabric controllers:

Management I/O Module Master LEDs		**Fabric Controller Master LEDs**		**Action**
In slot 15	**In slot 16**	**In slot 11**	**In slot 12**	
On	On	On	On	Too many master controllers are active. 1. Reboot the switch chassis. See "Rebooting the entire switch chassis" on page 145. 2. Replace the Fabric Controller Core in slot 11. 3. Replace the Fabric Controller Core in slot 12.
Off	Off	Off	Off	No master controller is active. 1. Reboot the switch chassis. See "Rebooting the entire switch chassis" on page 145 2. Replace the Fabric Controller Core in slot 11. 3. Replace the Fabric Controller Core in slot 12.
On	Not installed or Off	On	Not installed or Off	This is good status for the fabric controller. Go to the next step. (step 4)
On	Not installed	Off	Not installed or Off	No master controller is active. 1. Reboot the switch chassis. See "Rebooting the entire switch chassis" on page 145 2. Replace the Fabric Controller Core in slot 11.
Off	Not installed	Off	Not installed or Off	1. Reboot the switch chassis. See "Rebooting the entire switch chassis" on page 145 2. If the problem still exists, replace the Fabric Controller Core in slot 11.
Not installed or Off	On	Not installed or Off	On	This is good status for the fabric controller. Go to the next step. (step 4)
Not installed	Off	Not installed or Off	Off	1. Reboot the switch chassis. See "Rebooting the entire switch chassis" on page 145. 2. Replace the Fabric controller in slot 12.

4. Check the LEDs on the fabric controller module(s) in slots 11 and 12 using the procedure in “Fabric controller LEDs” on page 130, and perform any prescribed actions.
5. Check the LEDs on the management I/O module(s) in slots 15 and 16 using the procedure in “Switch management I/O module LEDs” on page 133, and perform any prescribed actions.
6. Replace the management I/O module(s) in the FRU list:
   - If this is a 7048-120 or SFS7000P, replace the entire switch chassis.
   - If this is a 7048-270 or SFS7008P, the I/O management modules are in slot 15 and slot 16. Replace the card in slot 15 first. If that does not fix the problem, replace the card in slot 16.
7. Try rebooting the HMC, in the event that the SNMP daemon stopped.
8. Try rebooting the switch using the CLI procedure in “Rebooting the entire switch chassis” on page 145. **This procedure ends here.**

IBNSSMM

No Master Subnet Manager was found on a subnet.

You are here because you have noted this isolation procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx.

Notes:

1. IBM service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.
2. After repair procedures, verify network function according to “Repair verification” on page 150.

Perform the following procedure:

1. Record the serviceable event reference code, serviceable event text, and FRU list.
2. Set a switch to master Subnet Manager using the procedure found in “Reordering Subnet Manager priority” on page 142.
3. Check for a master Subnet Manager using the procedure found in “Checking the subnet manager” on page 140
4. If you still do not have a master Subnet Manager, reboot the switch that you wish to be the subnet manager using the procedure in “Rebooting the entire switch chassis” on page 145
5. Call your next level of support and supply them with the above recorded information.
6. **This ends the procedure.**

IBNSSMR

The initial database synchronization, which is established between the master subnet manager and the highest-priority standby subnet manager, failed.

You should be here because you have noted this isolation procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx.

Note:

- IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service is responsible.
- For instructions on FRU replacement, refer to the appropriate service documentation.

- After repair procedures, verify network function according to the “Repair verification” on page 150 procedure.

1. Ensure the subnet manager on the node specified in the FRU list is operating correctly. For instructions on how to check on the subnet manager, refer to “Checking the subnet manager” on page 140.
2. If the subnet manager on the node specified in the FRU list is not operating correctly, restart the subnet manager; For instructions on restarting the subnet manager, refer to “Restarting the subnet manager” on page 142. Otherwise, adjust the cold synchronization timeout to a larger value. For instructions on adjusting the cold synchronization timeout value, refer to “Adjusting database-synchronization timeout” on page 161.
3. **This ends the procedure.**

IBNSSMU

There is a mismatch in the level of InfiniBand subnet managers used in this system.

You are here because you have noted this isolation procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx.

Note:

- IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service is responsible.
- After repair procedures, verify network function according to the “Repair verification” on page 150 procedure.

There is a mismatch in the level of InfiniBand subnet managers in this system. While this can be acceptable between successive versions of subnet managers, if the versions are too far apart, the likelihood of their functioning correctly is reduced.

1. Update the InfiniBand subnet managers so that all of them have identical software versions. See “Checking switch software levels” on page 144 and “Updating switch software” on page 144.
2. **This ends the procedure.**

IBNSSWE

A software error has occurred in interpreting an event. The fault could be with the subnet manager or the IBM Network Manager code.

You are here because you have noted this isolation procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx.

Note:

- IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service is responsible.
- After repair procedures, verify network function according to the “Repair verification” on page 150 procedure.

1. Record the serviceable event reference code, serviceable event text, and FRU list.
2. Call the next level of software support and supply them with the above recorded information.
3. **This ends the procedure.**

IBNSSWP

A software-initiated reboot has occurred.

You should be here because you have noted this isolation procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx.

Note: IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service is responsible.

Note: After repair procedures, verify network function according to the "Repair verification" on page 150 procedure.

Do the following:

1. Examine the ts_log file for detailed information about which application stopped and the circumstances for this event. Refer to "Reviewing the ts_log file" on page 159.
2. File a bug report. Refer to "Filing bug reports" on page 163.
3. If the problem occurs repeatedly, the user can either reconfigure the system, or upgrade the system software.
4. **This ends the procedure.**

IBNSTHM

Isolation procedure for a thermal event.

You should be here because you have noted this isolation procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx.

Note:

- IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service is responsible.
- For instructions on FRU replacement, refer to the appropriate manual using the references in "InfiniBand switch reference information" on page 3.
- After repair procedures, verify network function according to the "Repair verification" on page 150 procedure.

A thermal condition was detected in an InfiniBand switch. Look for all reports of thermal conditions, and do the following:

- If a critical condition exists:
  1. Ensure that proper airflow is being provided to the switch in the FRU list.
  2. Check the fans for the switch in the FRU list, and replace the fans if they are not operating correctly.
  3. Replace the FRUs in the FRU list according to the order in which they are presented.
  4. **This ends the procedure.**
- If a warning condition exists:
  1. Record the duplicate count in the Serviceable event detailed attributes panel in the Manage Serviceable Events – Serviceable Event Details window.
  2. Record the location code of the switch in the FRU list.
  3. Ensure that proper airflow is being provided to the switch in the FRU list.

4. Check the fans for the switch in the FRU list, and replace the fans if they are not operating correctly.
5. Perform the following procedure:
   a. Go to IBM Network Manager and open the View Management Properties window.
   b. Click the **Switch** tab.
   c. Select the switch chassis that contains the FRU from the FRU list.
   d. Click **Environmentals**.
   e. If the status is Normal and the duplicate count is fewer than 5, close the serviceable event, and end this procedure.
   f. If the duplicate count is greater than or equal to 5, or if the status is not Normal, replace the FRU in the FRU list.
   g. **This ends the procedure**.
6. Check other serviceable events for a normal condition being reported. If a normal condition is reported, consider the following
   a. If there are more warnings than normal conditions, consider replacing the switch as a preventive maintenance action.
   b. If there are an equal number of normal conditions versus warning conditions, the situation may have corrected itself.

      The environmental status should be Normal. If it is not, replace the FRU.
7. The change to a normal condition is also reported to Service Focal Point. Therefore, after performing any service actions against Critical or Warning conditions, look for a serviceable event normal condition (CB100002), and close it out.
8. **This ends the procedure.**

IBNSVPD

Use this procedure when the IBM Network Manager had difficulty reading VPD for other cards.

You should be here because you have noted this Isolation Procedure as the first FRU in the FRU list for a serviceable event that begins with CBxxxxxx.

Note:

- IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service is responsible.
- After repair procedures, verify network function according to the “Repair verification” on page 150 procedure.

Perform the following procedure:

1. Look for other similar serviceable events on the HMC running the IBM Network Manager.
2. If the IBM Network Manager had difficulty reading VPD for other cards, perform the following procedure. Otherwise, continue to the next step. It is likely that there is a single common failure point that is causing the problem.
   a. Inspect all Ethernet connections and ports in the service network to verify that they are operational. If you find a problem with the service network, take the appropriate action to fix it, and stop executing this procedure.

      If you do not find a problem with the service network, continue with the following step.
   b. Replace the management I/O module(s): If this is a 7048-120, replace the entire switch chassis.

      If this is a 7048-270, the I/O management modules are in slot 15 and slot 16. Replace the card in slot 15 first. If that does not fix the problem, replace the card in slot 16.

3. Replace the card in the FRU list.
4. **This procedure ends here.**

Finding the other side of the link

Use this procedure to find the other side of the link to a known device.

To find the other side of the link to a known device, use any of the following methods:

- Refer to cable planning documentation.
- Make a note of any cable labeling that indicates connectivity.
- Trace the cable from one end to the other.
- Use the IBM Network Managerto perform one of the following:
  - If the known device is a switch port, perform the following procedure:
    1. Open the Switch Topology View window.
    2. Find and select the known switch port in the window.
    3. Record the neighbor location code.
  - If the known attached device is a GX Host Channel Adapter (HCA), perform the following procedure:
    1. Open the End-point View window.
    2. Find and select the known attached GX HCA (adapter).
    3. Record the neighbor location code.

Interpreting LEDs

The LED interpretation covered by this section is referenced in the following table. The FRU LED column value is either Yes or No, and refers to whether the LED indicates whether a FRU is an error condition. If a particular FRU LED is available only on a particular switch model, that model is shown in parentheses.

LEDs	FRU LED?	Procedure Reference
Port	Yes	"7048-120 or SFS7000P switch port LEDs" on page 123
	Yes	"7048-270 or SFS7008P Switch port LEDs" on page 123
Entire Chassis Status	Yes	"Switch chassis and system-wide LEDs" on page 124
Power Supply status	Yes	"7048-270 or SFS7008P switch power supply LEDs" on page 126
Fan Tray (7048-270 and SFS7008P)	Yes	"Switch fan tray LEDs" on page 127
Line Interface Module (7048-270 and SFS7008P)	Yes	"Line Interface Module status indicator" on page 128
Fabric Controller Module	Yes	"Fabric controller LEDs" on page 130
Management I/O Module (7048-270 and SFS7008P) **Note:** The Management I/O Module contains LEDs for the module, but also LEDs that represent some LEDs that are seen only on the front of the chassis.	Yes	"Switch management I/O module LEDs" on page 133
Master Management Connection (7048-270 and SFS7008P)	No	"Switch master management connection LED" on page 134
Ethernet Management Port (7048-270 and SFS7008P)	No	"Ethernet management port LEDs" on page 135
GX bus adapter status	No	"GX bus adapter LEDs" on page 136

Note: For complete information about Topspin switch LEDs, see "InfiniBand switch reference information" on page 3.

7048-120 or SFS7000P switch port LEDs

Use the following table to interpret the switch port LED states. If the LED Color (state) is Off, proceed to the procedure following the table.

LED color (state)	Description
Off	Logical link has not been established. See the following Port LED is Off procedure.
Green (on continuously)	Logical link established. This is a normal condition.
Green (blinking on and off)	Logical link with activity. This is a normal condition.

Port LED is Off: If the state of the LED is Off, perform the following steps in sequence, proceed to the next step only if the current step does not fix the problem. When the LED changes to a normal state, the problem fixed.

1. If the LED on the other side of the link indicates a normal condition, the link is most likely operating correctly. The problem is probably in the LED. Replace this switch card during the next maintenance opportunity.
2. Verify that the cable is plugged in and seated at both ends. Fix any other obvious problems that you might find.
3. Check for bent pins on the cable. If bent pins exist, replace the cable.
4. Run link diagnostics to isolate the problem. Perform any service actions recommended in the procedure found in "Link diagnostic procedures" on page 165.
5. Replace the cable.
6. If there is an adapter on this link, replace the adapter.

   Notes:

   a. Any subsequent repair action recommended in this procedure might cause multiple link outages that can negatively impact customer performance. Consider deferring maintenance to a time that will minimize impact to customer performance.
   b. Before replacing an HCA, review "Installing or replacing a GX Host Channel Adapter" on page 64. If you are considering deferred maintenance of the adapter, review "Deferring replacement of a failing Host Channel Adapter" on page 66.
7. Replace the switch card with the LED that is off.
8. If the other side of the link is a switch, replace the switch card on the other side of the link.
9. **This ends the procedure.**

7048-270 or SFS7008P Switch port LEDs

There are two LEDs per switch port on a 7048-270 or SFS7008P.

Use the following table to interpret switch port LED states. If the state of the LED is off, proceed to the procedure following the table.

Table 21.

Top (tx/rx) LED	Bottom (link) LED	Indication
Off	Off	No link signal is detected. See the procedure below.

Table 21. (continued)

Top (tx/rx) LED	Bottom (link) LED	Indication
Off	Solid	Physical link is established. No logical link is established. If this link supports a 12x connection with the use of an octopus cable, this may be normal. Refer to "Behaviors of octopus cable connections in static-12x configurations" on page 19. Otherwise, use the procedure below.
Solid	Solid	Physical and Logical link established. This is a normal condition
Blinking	Solid	Physical and Logical link are established and there is traffic activity. This is a normal condition

If both LEDs are off, or only the Top (tx/rx) LED is off, and the cable connected to it is not an octopus cable:

Perform the steps in sequence proceeding to the next step only if the current step did not fix the problem. When the LEDs go to a normal state, consider the problem fixed.

Note: If there are no performance problems being reported, consider performing this procedure and any FRU replacements during the next maintenance window.

1. If the LED on the other side of the link indicates a normal condition, the link is operating correctly and there is something wrong with the LEDs on the FRU that has both LED off. Replace this LIM (Line Interface Module) during the next maintenance window. **This procedures ends here.**
2. Verify that the cable is plugged and seated at both ends. Fix any problems that you might find.
3. Check for bent pins on the cable. If it has bent pins, replace the cable.
4. Run link diagnostics to isolate the problem. Perform any service actions recommended in the procedure found in "Link diagnostic procedures" on page 165.
5. Replace the cable.
6. If there is an adapter on this link, replace the adapter.

   Notes:

   a. Any subsequent repair action recommended in this procedure might cause multiple link outages that can negatively impact customer performance. Consider deferring maintenance to a time that will minimize impact to customer performance.

   b. Before replacing an HCA, review "Installing or replacing a GX Host Channel Adapter" on page 64. If you are considering deferred maintenance of the adapter, review "Deferring replacement of a failing Host Channel Adapter" on page 66.
7. Replace the switch card that has its LED in the off state. If the other side of the link is a switch, replace the switch card on the other side of the link.

This procedure ends here.

Switch chassis and system-wide LEDs

The chassis LED on the 7048-120 and SFS7000P is the same indicator as the system-wide LED on the 7048-270 and SFS7008P.

Color (state)	Description
Off	No system power or LED failure.
Yellow (solid)	Operator intervention is required. A system error was detected, such as a fan error, a Power On Self Test (POST) failure, or a power supply failure. This will only be on while the error is still present. If the error is transient, the color of the LED may change when the error is no longer present.

Color (state)	Description
Yellow (blinking)	Initiated automatically during the LED test that follows the application of power (16 seconds). This is also used by the IBM Network Manager as an identification LED for the chassis.
Solid Green	Indicates correct operation and that there are no critical errors.

Chassis LED is off

If the chassis LED is off and you believe it should be on, perform the following procedure:

1. Determine if the LED is faulty:
   a. In the IBM Network Manager, open the Switch Topology view.
   b. If the switch power is on, there is a problem with the LED. You may choose to defer maintenance, because this is a disruptive repair action:
      - For the 7048-120 or SFS7000P, replace the entire switch chassis.
      - For the 7048-270 or SFS7008P, replace the Chassis ID card in slot 17.
2. Verify that power cables are plugged in and that the power supplies are correctly seated.
3. If a power supply is not correctly seated, reseat it.
4. Replace the power supplies one at a time.
5. **This procedure ends here.**

Chassis LED is yellow (solid)

If the chassis LED is yellow and solidly on, there is an error condition on this switch. Perform the following procedure:

Note: Whenever you replace a cable, line interface module, fabric controller module, or an adapter, you will disrupt service to one or more links. Expect performance degradation and the potential for serviceable events to be generated as a result of the service action. If the service action will disrupt service to one or more links, and the customer is not currently experiencing degraded performance, you can defer maintenance.

If there is an adapter or switch that was powered off and connected to this switch, this is the likely cause of the error. However, use the following procedure to verify that this is the case.

1. Go to Service Focal Point on the HMC that is running the IBM Network Manager, and look for a serviceable event related to this switch or any device connected to this switch.
2. If there are no serviceable events on the HMC that is running the IBM Network Manager, service any events on all other HMCs that are reported by adapters or system units that are connected to this switch.
3. If none of the previous actions have resolved the problem, examine the FRU LEDs in the switch chassis, and replace any FRUs that indicate a problem. For the list of FRU LEDs see “Interpreting LEDs” on page 122.
4. Call your next level of support.
5. **This procedure ends here.**

Chassis LED is yellow (blinking)

If the chassis LED is yellow and blinking, this indicates a test state. Normally, this occurs when the switch goes through a power-on or reboot sequence. The IBM Network Manager can also put the LED in this state for chassis-identification purposes.

If you believe that the chassis LED should not be yellow and blinking, perform the following procedure:

1. Open the Switch Topology View window in the IBM Network Manager.
2. Find and select the switch chassis that has the LED that is yellow and blinking.
3. From the menu, click **Selected > Properties**.
4. Click the **System** tab.
5. Check the **Identify LED** property.
   a. If the LED is on and it should be off, do the following:
      1) Go to the IBM Network Manager and open the Switch Topology view.
      2) Select the chassis that has the flashing LED.
      3) From the menu, click **Selected > Identify > Off**.
   b. If the LED is off, proceed to the next step.
6. Check the LED again to be sure that the switch was not going through the portion of the boot sequence that tests the chassis LED.
7. If the chassis LED is still yellow and blinking, reboot the switch. You must close out the serviceable events that are caused by the switch reboot.
8. If the chassis LED is still yellow and blinking, your next action depends on the switch model:

   Note: Whenever you replace a switch chassis, you will disrupt service to one or more links. Expect performance degradation and the potential for serviceable events to be generated as a result of the service action. If the links appear usable, can defer maintenance to a more convenient time for the customer.

   - If the switch is a 7048-120 or SFS7000P, replace the switch chassis.
   - If the switch is a 7048-270 or SFS7008P, replace the Chassis ID module in slot 17.

7048-270 or SFS7008P switch power supply LEDs

The following table describes the various power supply LEDs:

LED color	Description
Off	DC output failure
Green (solid)	AC connected, DC output is acceptable
Yellow (off)	No failure on the power supply
Yellow (solid)	Operator intervention is required. A failure was detected within the power supply.
Yellow (blinking)	Identify. When the LED is in this state, it helps to identify a particular field replaceable unit on the chassis. This state must be initiated by the user. This LED can be initiated manually by using the **diag rack-locator** command in the global configuration mode. For more information, see the *Command Line Interface Reference Guide,* order number: 10-00012-07-A0.

Note: Power supply modules are redundant. However, after you have lost a power supply, redundancy no longer exists. Repairing the power supply module should not disrupt the network. Therefore, deferred maintenance is only warranted when you do not have a replacement part on hand.

Switch power LED is off

If the switch power LED is off, there is a DC output failure. Perform the following procedure:

1. Check the seating of the power cord for the power supply.
2. Check the seating of the power supply in the chassis.
3. If the power supply is incorrectly seated, reseat it.
4. Replace the power supply.

5. **This procedure ends here.**

Switch power LED is yellow (solid)

This indicates that an error is occurring within the power supply. Perform the following procedure:

1. Go to Service Focal Point on the HMC that is running the IBM Network Manager, and look for a serviceable event related to this switch or any device connected to this switch.
2. If the switch Power LED is still yellow (solid), replace the power supply.
3. Call your next level of support.
4. **This procedure ends here.**

Switch power LED is yellow (blinking)

The yellow blinking LED is used to identify the switch power supply. If this LED is blinking and you believe that it should not be blinking, perform the following procedure:

1. Record which switch power module has the blinking LED.
2. Wait a few minutes for the yellow LED to stop blinking, because it is possible that the switch is going through a reboot.
3. If the LED is still blinking, because the IBM Network Manager does not provide a method for turning on or off the switch power FRU identification LED, you must use the Topspin administration tasks to turn it off. See "Accessing FRU Identification LEDs" on page 162.
4. **This procedure ends here.**

Switch fan tray LEDs

The switch fan tray is found only in a 7048-270 or SFS7008P switch, because the fans are integrated with the power modules in the 7048-120 or SFS7000P.

LED color	Description
Green and Yellow off	No power to the fan tray, or an LED failure
Green (solid)	Fans in the fan tray are running with no detected errors
Green (off)	No power to the fan tray, LED failure, or yellow LED is on.
Yellow (off)	Fan tray running with no errors detected.
Yellow (solid)	Operator intervention is required. A failure was detected within the fan tray.
Yellow (blinking)	Assists in identifying a particular field replaceable unit on the chassis.

Note: Not all of the preceding states have a specific procedure. Some are contained within another overall procedure.

Fan trays are redundant. However, after a fan tray fails, redundancy no longer exists. Repairing the fan tray should not disrupt the network. Therefore, deferred maintenance is only warranted when you do not have a replacement part on hand.

Switch fan tray LED is green (off)

If the green LED is off, it can indicate that there is either no power to the fan tray, or that there is a fan error. Do one of the following to isolate the condition:

- If the yellow LED is blinking, go to "Switch fan tray is yellow (blinking)" on page 128.
- If the yellow LED is on solidly, go to "Switch fan tray LED is yellow (on)" on page 128.
- If the yellow LED is off, this indicates that there is no power applied to the fan tray. If this is unexpected (that is, the switch is powered on), perform the following procedure:

1. Verify that the power cables are attached and plugged into an active power source.
2. Verify that the power supplies are appropriately powered; see "7048-270 or SFS7008P switch power supply LEDs" on page 126. Perform any recommended service in that procedure.
3. Verify the seating of the fan tray.
4. If the fan tray is incorrectly seated, reseat it.
5. If both LEDs are still off, replace the fan tray.
6. **This ends the procedure.**

Switch fan tray LED is yellow (on)

This condition indicates that there is an error being reported by the fan tray. Perform the following procedure:

1. Go to Service Focal Point on the HMC that is running the IBM Network Manager, and look for a serviceable event related to this switch or any device connected to this switch.
2. If the switch Fan LED is still yellow (solid), replace the fan tray.
3. Call your next level of support.
4. **This procedure ends here.**

Switch fan tray is yellow (blinking)

This condition identifies the Switch Fan Tray FRU. If the LED is blinking, and you believe that it should not be blinking, perform the following procedure:

1. Record which fan tray has the blinking LED.
2. Wait a few minutes for the yellow LED to stop blinking, because it is possible that the switch is going through a reboot.
3. If the LED is still blinking, you must access the Topspin administration tasks to turn it off. This is because the IBM Network Manager does not provide a method for turning the switch fan tray FRU identification LED on or off. See the procedure in "Accessing FRU Identification LEDs" on page 162.
4. **This procedure ends here.**

Line Interface Module status indicator

The Line Interface Module status-indicator LEDs are available only on the 7048-270 or SFS7008P. In addition to the LEDs for each port, the Line Interface Module (LIM) status indicator provides the status across the entire LIM. The LIM status indicator LEDs are located to the far left of the LIM. Use the following table to identify the correct procedure to follow for each state.

LED color	Description
Green and Yellow off	No power to the LIM, or an LED failure
Green (solid)	The LIM is running with no detected errors
Green (off)	No power to the LIM, LED failure, or yellow LED is on.
Yellow (off)	No errors detected on the LIM
Yellow (solid)	Operator intervention required. Failure detected within the LIM. The "!" label indicates a failure.
Yellow (blinking)	Assists in identifying a particular field replaceable unit on the chassis.

Note: Not all states in the preceding table have a specific procedure.

A key step in each procedure for the LIM status indicator LED is to check all the port LEDs on the LIM. This may help you understand the scope of the condition.

If you must replace a LIM, all of the ports connected to that LIM will be unavailable until the new LIM is installed. If there are functional ports on the LIM, you can defer maintenance to a time that will minimize impact to customer performance.

Switch LIM status indicator LED is green (off)

If the green LED is off, it can indicate that there is either no power to the LIM, or that there is an error on the LIM. Do one of the following to isolate the condition:

- If the yellow LED is blinking, go to "Switch LIM status indicator LED is yellow (blinking)" on page 130.
- If the yellow LED is on solidly, go to "Switch LIM status indicator LED is yellow (on)."
- If the yellow LED is also off, this indicates that there is no power applied to the LIM. If this is unexpected (that is, the switch is powered on), perform the following procedure:
  1. Verify that the power cables are attached and plugged into an active power source.
  2. Verify that the power supplies are properly powered; go to "7048-270 or SFS7008P switch power supply LEDs" on page 126. Perform any recommended service in that procedure.
  3. Verify the seating of the LIM.
  4. If the LIM is incorrectly seated, reseat it.
  5. If both the yellow and green LEDs are still off, replace the LIM. **This procedure ends here.**

Switch LIM status indicator LED is yellow (on)

This condition indicates that there is an error being reported by the switch LIM. Perform the following procedure:

1. Go to Service Focal Point on the HMC that is running the IBM Network Manager, and look for a serviceable event related to this switch or any device connected to this switch.
2. If the LIM LED is still yellow (solid), check the LEDs on each of the ports on the LIM; go to "7048-120 or SFS7000P switch port LEDs" on page 123.
   - If only a single port appears to be faulty (green LED off), do the following procedure:
     a. Run diagnostics against that link to isolate the fault; go to "Link diagnostic procedures" on page 165. Perform any recommended service actions based on diagnostics results.
     b. If diagnostics do not reveal a faulty FRU, replace the LIM.
   - If more than one port appears faulty, perform the following procedure:

     Note: The following procedure renders the LIM ports nonfunctional. If there are functional ports on the LIM, and the customer is satisfied with current performance, consider deferring this procedure until a time that will minimize impact to customer performance.

     a. Make note of the slot in which the LIM is located.
     b. Run card diagnostics, and perform any recommended service actions; go to "InfiniBand switch card diagnostics" on page 173.
     c. If the switch card diagnostics pass, and you have performed this procedure previously, this indicates a persistent fault that cannot be isolated through diagnostics. In that case, replace the LIM.
     d. If this is the first pass through this procedure, continue on to the next step, but make a note that you have run this procedure against this LIM.
     e. If, after re-enabling the card, the yellow light is still on solidly, replace the LIM. If there are functional ports on the LIM and the customer is satisfied with the current level of performance, consider deferring maintenance until a time that will minimize impact to customer performance.
3. **This procedure ends here.**

Switch LIM status indicator LED is yellow (blinking)

The Switch LIM status indicator LED identifies the Switch LIM FRU. If it is blinking and you believe that it should not be blinking, perform the following procedure:

1. Record the slot in which the LIM is plugged.
2. Wait a few minutes for the yellow LED to stop blinking, because it is possible that the switch is going through a reboot.
3. If the LED is still blinking, go to "Manipulating FRU identification LEDs" on page 136 to turn off the LED.
4. **This procedure ends here.**

Fabric controller LEDs

Fabric controller LEDs apply only to the 7048-270 or SFS7008P. The Fabric Controller module status LEDs indicate status of the module. The active system master indicator indicates when the fabric controller is the master module for the chassis. The following describes the fabric controller LEDs.

Fabric controller module status	
LED color	**Description**
Green and Yellow off	No power to the fabric controller, or an LED failure
Green (solid)	The fabric controller is running with no detected errors
Green (off)	No power to the fabric controller LED, LED failure, or yellow LED is ON.
Yellow (off)	There are no errors detected on the fabric controller.
Yellow (solid)	Operator intervention required. Failure detected within the fabric controller. The "!" label indicates a failure.
Yellow (blinking)	Assists in identifying a particular Field Replaceable Unit on the chassis.

Active system master indicator for the fabric controller	
LED color	**Description**
Green on	This module is currently the system master. Only fabric modules in one of the two core slots can become master.
Green (off)	This module is not the active system master. Only fabric modules in one of the two core slots can become master.

Note:

- Not all of the preceding states have a specific procedure. Some are contained within another overall procedure.
- If you must replace a fabric controller, this may severely impact customer performance. Therefore, you can defer maintenance until a time that will minimize impact to customer performance.

Fabric controller module status indicator LED is green (off)

If the green LED is off, it can indicate that there is either no power to the fabric controller, or that there is an error in the fabric controller. Use the following procedure to isolate the condition:

- If the yellow LED is blinking, go to "Fabric controller module status indicator LED is yellow (blinking)" on page 131.
- If the yellow LED is on solidly, go to "Fabric controller module status indicator LED is yellow (on)" on page 131.

- If the yellow LED is also off, this indicates that there is no power applied to the fabric controller. If this is unexpected (that is, the switch is powered on), perform the following procedure:
  1. Verify that the power cables are attached and plugged into an active power source.
  2. Verify that the power supplies are correctly powered; go to "7048-270 or SFS7008P switch power supply LEDs" on page 126. Perform any recommended service in that procedure.
  3. Verify the seating of the fabric controller.
  4. If the fabric controller is incorrectly seated, reseat it.
  5. If both the yellow and green LEDs are still off, replace the fabric controller. **This procedure ends here.**

Fabric controller module status indicator LED is yellow (on)

This indicates that there is an error being reported by the switch LIM. Perform the following procedure:

1. Go to Service Focal Point on the HMC that is running the IBM Network Manager, and look for a serviceable event related to this switch or any device connected to this switch.
2. If the Fabric Controller LED is still yellow (solid), perform the following procedure:

   Note: The following procedure renders the fabric controller nonfunctional, which may affect many ports on the switch. If the customer is satisfied with current performance, consider deferring this procedure until a time that will minimize impact to customer performance.

   a. Make note of the slot in which the fabric controller is located.
   b. Run card diagnostics, and perform any recommended service actions; see "InfiniBand switch card diagnostics" on page 173.
   c. If the switch card diagnostics pass, and you have performed this procedure previously, this indicates a persistent fault that cannot be isolated through diagnostics. In that case, replace the fabric controller.
   d. If this is the first pass through this procedure, continue on to the next step, but make a note that you have run this procedure against this fabric controller.
   e. If, after re-enabling the card, the yellow light is still on solidly, replace the fabric controller. **This procedure ends here.**

Fabric controller module status indicator LED is yellow (blinking)

This is used to identify the Switch Fabric Controller FRU. If it is blinking and you believe that it should not be blinking, perform the following procedure:

1. Record the slot in which the fabric controller is plugged.
2. Wait a few minutes for the yellow LED to stop blinking, because it is possible that the switch is going through a reboot.
3. If the LED is still blinking, you must perform the following procedure to stop it from blinking: see "Manipulating FRU identification LEDs" on page 136.
4. **This procedure ends here.**

Fabric controller active system master indicator is green (off)

When the fabric controller active system master indicator LED is off, this indicates that this card is not the active master. This LED is only valid for fabric controllers in one of the two core slots (switch card slot 11 and slot 12); see "Switch FRU Locations for the 7048-270 or SFS7008P" on page 220.

If you believe that this card should be the master controller card, perform the following procedure:

Note: This procedure could interrupt functionality on this switch and have a negative impact on customer performance. If the customer is satisfied with the current performance, consider deferring maintenance until a time that has minimal impact on customer performance.

1. Record the card slot for this card.
2. If this card is in a node slot (slots 9, 10, 13 and 14), it is valid that the active system master indicator LED is not on; therefore, this procedure ends here. Otherwise, proceed to the next step.
3. Verify that no other fabric controller card in a core slot is master by checking the other active system master indicator LEDs. Core slots are slot 11 and slot 12.
4. If the other fabric controller in a core slot is not the master assigned to this switch, perform the following procedure:
   a. Go to Service Focal Point on the HMC that is running the IBM Network Manager, and perform any service recommended on this switch.
   b. If all active system master indicators are still off, continue to the next step. Otherwise, this procedure ends here.
   c. Run card diagnostics on this fabric controller card; see "InfiniBand switch card diagnostics" on page 173. Perform any recommended service actions.
   d. If the card passes self-test diagnostics, run card diagnostics on the other fabric controller in a core slot (slot 11 or slot 12). Perform any recommended service actions.
   e. Check the active master indicator LED on both fabric controllers in core slots (slot 11 an slot 12).
   f. If there is still no fabric controller that is indicated to be the master, reboot the switch by performing the procedure found in "Rebooting switch cards" on page 160. In this instance, use the all parameter when choosing the slot to reboot. Perform any recommended service actions.
   g. If there is still no active system master LED on for any of the two fabric controllers in core slots in this switch chassis, begin replacing the fabric controllers one at a time until one of them is indicating that it is master. Begin with slot 11. Then, proceed to slot 12.
   h. If there is still no Active System Master LED on for any of the two Fabric Controllers in core slots in this switch chassis, replace the system backplane.
5. **This procedure ends here.**

Fabric controller active system master indicator is green (on)

When the active system master indicator LED is green and on, it indicates that this card believes that it is the master for the switch chassis in which it is located. Only one fabric controller in a chassis should be assigned as the master. If both the core slots fabric controllers (slot 11 and slot 12) are indicating that they are master, perform the following procedure:

Note: This procedure could interrupt functionality on this switch and have a negative impact on customer performance. If the customer is satisfied with his current performance, you may wish to consider deferring maintenance to a time that would have minimal impact on customer performance

1. Record the switch card slot into which this fabric controller is plugged.
2. If this fabric controller is in a node slot (slots 9, 10, 13 and 14), the active master indicator LED should not be on. Replace the fabric controller card. Otherwise, proceed to the next step.
3. Go to Service Focal Point on the HMC that is running the IBM Network Manager, and perform any recommended service on this switch.
4. If both active system master indicators are still on, continue to the next step. Otherwise, this procedure ends here.
5. Run card diagnostics on this fabric controller card; see "InfiniBand switch card diagnostics" on page 173. Perform any recommended service actions.
6. If the card passes self-test diagnostics, run card diagnostics on the other fabric controller in a core slot (slot 11 or slot 12). Perform any recommended service actions.

7. Check the active master indicator LED on both fabric controllers in core slots (slot 11 an slot 12).
8. If both active system master indicators are still on, reboot the switch by performing the procedure found in “Rebooting switch cards” on page 160. In this instance, you should use the all parameter when choosing the slot to reboot. Perform any recommended service actions.
9. If both active system master indicators are still on, begin replacing the fabric controllers one at a time until one of them is indicating that it is master. Begin with slot 11. Then, proceed to slot 12.
10. If both active system master indicators are still on, replace the system backplane.
11. **This procedure ends here.**

Switch management I/O module LEDs

This section applies only to the 7048-270 or SFS7008P switches.

The switch management I/O module LEDs reflect the states of LEDs seen on the front of the 7048-270 or SFS7008P, and other LEDs that apply to the switch management only.

- System-Wide LED; see “Switch chassis and system-wide LEDs” on page 124.
- Management Module Status; see the following table
- Power Supply 1 and 2 status; see “7048-270 or SFS7008P switch power supply LEDs” on page 126.
- Fan Tray 1 and 2 status; see “Switch fan tray LEDs” on page 127.
- Ethernet Management port status; see “Ethernet management port LEDs” on page 135.
- Master Management Connection LED; see “Ethernet management port LEDs” on page 135.

LED color	Description
Green and Yellow off	No module power, or LED failure.
Green (solid)	The Management Interface module is running with no errors detected.
Green (off)	No power to the Management Interface module, LED failure, or yellow LED is ON.
Yellow (off)	Fan tray running with no errors detected.
Yellow (solid)	Failure of the Management Interface module. Operator attention required.
Yellow (blinking)	Identify a Management Interface module. Assists in identifying a particular Field Replaceable Unit on the chassis.

Switch management I/O module LED is green (off)

The green LED being off can indicate that there is either no power to the fan tray, or that there is a fan error. Use the following procedure to isolate the condition:

1. If the yellow LED is blinking, go to “Switch management I/O module is yellow (blinking)” on page 134.
2. If the yellow LED is on solidly, go to “Switch management I/O module LED is yellow (on)” on page 134.
3. If the yellow LED is also off, this implies that there is no power applied to the fan tray. If this is unexpected (for example, the switch is powered on), perform the following procedure:
   a. Verify that the power cables are attached and plugged into an active power source.
   b. Verify that the power supplies are correctly powered; see “7048-270 or SFS7008P switch power supply LEDs” on page 126.
   c. Verify the seating of the management I/O module.
   d. If both LEDs are still off, replace the Management I/O module.
4. **This procedure ends here.**

Switch management I/O module LED is yellow (on)

This indicates that there is an error being reported by the fan tray. Perform the following procedure:

1. Go to Service Focal Point on the HMC that is running the IBM Network Manager and look for a serviceable event related to this switch or any device connected to this switch.
2. If the switch management I/O module LED is still yellow (solid), replace the switch management I/O module.
3. Call your next level of support.
4. **This procedure ends here.**

Switch management I/O module is yellow (blinking)

This is used to identify the Switch Fan Tray FRU. If it is blinking and you believe that it should not be blinking, perform the following procedure:

1. Wait a few minutes for the yellow LED to stop blinking, because it is possible that the switch is going through a reboot.
2. If the LED is still blinking, because the IBM Network Manager does not provide a method for turning on or off the switch power FRU identification LED, you must access the Topspin administration tasks to turn the LED off. See the procedure in "Accessing FRU Identification LEDs" on page 162.
3. **This procedure ends here.**

Switch master management connection LED

The following switch master management connection LED information only applies to a 7048-270 or SFS7008P.

LED color	Description
Green (off)	The management module is not connected to the active master controller, or is connected to the system's standby master
Green (solid)	The management module is connected to the system's active master controller

Switch master management connection LED is green (off)

When the Switch Master Management Connection LED is Green (off), this condition indicates that this Management I/O Module is not connected to the active master controller, or is connected to the system's standby master. If you believe this may be a problem, perform the following procedure:

1. Check the other Management I/O Module's Switch Master Management Connection LED. If it is on, this procedure ends here. Otherwise, proceed to the next step.
2. Verify that at least one of the Fabric Controllers in a core slot (slot 11 and slot 12) has its Active System Master LED on; see "Fabric controller LEDs" on page 130. If this is not the case, perform the recommended service action(s) in the referenced procedure. Otherwise, proceed to the next step.
3. Run switch card diagnostics against each Management I/O module (card slots 15 and 16), and perform any recommended service actions. See "InfiniBand switch card diagnostics" on page 173.
4. If it is still true that neither Management I/O Module's Switch Master Management Connection LED is Green (solid), begin replacing FRUs in the following order:
   a. Management I/O module in slot 15
   b. Management I/O module in slot 16
   c. Fabric Controller in slot 11
   d. Fabric Controller in slot 12
   e. backplane

5. **This procedure ends here.**

Switch master management connection LED is green (solid)

When the Switch Master Management Connection LED is Green (off), this indicates that this Management I/O Module is connected to the active master controller. This should be a normal condition. Because the Management I/O modules are paired with a fabric controller, only one Management I/O module should have Green (solid) Switch Management Connection LED on at any given time. If both Management I/O Modules' Switch Master Management Connection LED is Green (solid), perform the following procedure:

1. Check the Fabric Controllers in the core slots (slot 11 and slot 12) to see if they both have their Active System Master Indicator LED on.
2. If both Active System Master Indicator LEDs are on, perform the recommended service actions in "Fabric controller active system master indicator is green (on)" on page 132. Otherwise, proceed to the next step.
3. Because only one Fabric Controller's Active System Master Indicator LED is on, there is likely a problem with one of the Management I/O cards. Run diagnostics against the Management I/O cards and perform any recommend service actions in "InfiniBand switch card diagnostics" on page 173.
4. If it is still true that both Management I/O Module's Switch Master Management Connection LED is Green (solid), perform the following procedure:
   a. If the Fabric Controller in slot 11 is currently indicating that it is the Master Management module, replace FRUs in the following order, until the problem is fixed.
      1) Management I/O module in slot 16
      2) Fabric Controller in slot 12
      3) Management I/O module in slot 15
      4) Fabric Controller in slot 11
   b. If the problem still exists, replace the backplane
5. **This procedure ends here.**

Ethernet management port LEDs

The following Ethernet management port LED information applies only to a 7048-270 or SFS7008P.

LED color	Description
Left port LED (off)	Port is not connected or logical link is down.
Left port LED (on)	Port is connected and logical link is up. This is a good condition.
Left port LED (on) and Right port LED (off)	No traffic is moving on the Ethernet Management port link. Logical link may or may not be up.
Right port LED (on)	Traffic is moving on the Ethernet Management port link. This is a good condition.

Ethernet management port (left) LED is off

The Ethernet management port LED indicates that the port is not connected or the logical link is down. Perform the following procedure:

1. Verify the cable connection to this port, and the port on the other end of the cable.
2. Verify that the port on the other side of the cable is up and operational.
3. If everything in the Ethernet network appears to be well seated and operational, replace the management I/O module.
4. **This procedure ends here.**

GX bus adapter LEDs

Refer to the appropriate guide for LED descriptions for the server in which the GX Bus adapter is installed. This information is available in the IBM Systems Hardware Information Center.

Manipulating FRU identification LEDs

If you need to turn on/off FRU identification LEDs, first determine if you can use IBM Network Manager's functions to do so, or if you need to use the switch's Command Line Interface (CLI).

Go to the procedure in the following table that is indicated by the LEDs that you want to manipulate.

Component LED	Description
Switch Chassis	"Using IBM Network Manager to access FRU identification LEDs"
Switch Line Interface Module (LIM)	"Using IBM Network Manager to access FRU identification LEDs"
Switch Port	"Using IBM Network Manager to access FRU identification LEDs"
Switch Fabric Controller	"Accessing FRU Identification LEDs" on page 162
Switch Power Module	"Accessing FRU Identification LEDs" on page 162
Switch Fan Tray	"Accessing FRU Identification LEDs" on page 162
Switch Management I/O	"Accessing FRU Identification LEDs" on page 162

Using IBM Network Manager to access FRU identification LEDs

1. Go to the Switch Topology View window in the IBM Network Manager.
2. Select the switch chassis that has the LED that is yellow and blinking.
3. From the menu, click **Selected > Properties**.
4. Click the **System** tab.
5. Check the Identify LED property.
   a. If the LED is On and it should be off, perform the following:
   b. If the LED is Off, proceed to the next step.
6. Check the LED again ensure that the switch was not going through the portion of the boot sequence that tests the chassis LED.
7. If the chassis LED is still yellow and blinking, reboot the switch.

   Close out the serviceable events that are caused by the switch reboot.
8. If the chassis LED is still yellow and blinking, your next action depends on the switch model:

   Note: Whenever you replace a switch chassis, service to one or more links is disrupted. Expect performance degradation and the potential for serviceable events to be generated as a result of the service action. If the links appear usable, you may wish to defer maintenance to a more convenient time for the customer.

   If the switch is a 7048-120, replace the switch chassis.

   If the switch is a 7048-270, replace the Chassis ID module in slot 17.

Isolating a problem with a port or link

This is a generic procedure for isolating problems with a port or a link.

To isolate a problem with a port or link, do the following:

1. Record the location code of the device(s) that has the port or the link. The information may be available from the following locations:
   - The FRU list for a serviceable event
   - An IBM Network Manager window
2. Repeat the following procedures for each device on the link.
   - If the IBM Network Manager is being used, you can get the other side of the link by following the procedures in "Finding the other side of the link" on page 122 and then viewing the properties on the port in question.
   - If the IBM Network Manager is not being used, determine the other side of the link using the Element Manager, see "InfiniBand switch reference information" on page 3 for more information about Element Manager or cable planning documentation for the cluster.
3. If there is no IBM Network Manager and the device is a switch, go to the Element Manager, see "InfiniBand switch reference information" on page 3, for more information on Element Manager, and look for errors reported by the device and perform the prescribed service procedures.
4. Go to Service Focal Point (SFP) on the HMC that is running the IBM Network Manager. If there is a serviceable event reported with the device's location code in the FRU list, perform the prescribed service procedures.
5. Go to SFP on the HMC that is controlling the server in which the device is populated. If there is a serviceable event reported with the device's location code in the FRU list, perform the prescribed service procedures.
6. If there are no serviceable events reported with the device in the FRU list, check the cable connections on the link to be sure that they are correctly seated.

   Note: If this is a 4x link configured to run at 12x with an octopus cable, verify the seating of all the cable connectors for the octopus cable. Be careful to record which 4x connector goes to which 4x switch port. If you do not do this, you need to determine the configuration again using the procedure in "Planning octopus cables in static 12x cabling" on page 19. For more on behavior of 12x connections see "Planning octopus cables in static 12x cabling" on page 19.
7. Verify that the LEDs on each port are in the correct state, and perform any service as described in "Interpreting LEDs" on page 122.

   Note: If this link has 4x switch ports configured to run at 12x with an octopus cable, for the LED behavior for groups of 4x switch ports connected to an octopus cable, refer to "Verifying static-12x mode connectivity" on page 67.
8. If one of the devices is an adapter, use the operating system to determine if the device is seen and available by the operating system. If it is not, perform the prescribed service procedure.
9. Run link diagnostics. See "Link diagnostic procedures" on page 165
10. Replace FRUs in the following order:
    a. The cable

       Note: If you are replacing an octopus cable, you will need to connect the switch ports first, and then the HCA side. Be sure to record which 4x connectors attach to which 4x switch ports. Then, verify the operation of the link using "Verifying static-12x mode connectivity" on page 67. For more information on 12x connections, see "Planning octopus cables in static 12x cabling" on page 19.
    b. The adapter (if one is on the link)

       Note: If you are replacing an HCA that has an octopus cable attached to it, you will need to verify the operation of the link using "Verifying static-12x mode connectivity" on page 67. If you have removed the switch port connectors, replace those before you connect the HCA side of the octopus cable. For more information on 12x connections see "Planning octopus cables in static 12x cabling" on page 19.

c. Switch card(s)

Note:

- If you are replacing a switch card that has octopus cables attached to it:
  - Configure all of the ports on the new switch cards that are acting as a group in a 12x link using an octopus cable. See "Configuring Static-12x groups" on page 20.
  - Verify their operation. See "Verifying static-12x mode connectivity" on page 67.

  For more information on 12x connections, see "Planning octopus cables in static 12x cabling" on page 19.

11. Call your next level of support.
12. **This procedure ends here.**

Isolating performance problems

This is a generic procedure for isolating performance problems.

Performance degradation can result from several different problems, including:

- a hardware failure
- Installation problems
- Configuration issues

Before calling your next level of service, do the following to isolate a performance problem:

1. Look for hardware problems:
   a. Open Service Focal point on all HMCs and perform prescribed service any open serviceable events. If you have redundant HMCs configured, you need only open Service Focal Point on one HMC in each set of redundant HMCs.

      Note: Recall that switch ports connected to PCI-X HCAs may report serviceable events that are induced by events in servers with PCI-X HCAs. Examples of such events are server reboots, checkstops, or adapter errors that cause link errors.

   b. Open the IBM Network Manager on the HMC that is running the IBM Network Manager. Open each of the following windows within the IBM Network Manager and perform any service tasks prescribed for the status found within those windows. Refer to "Status procedures for the IBM Network Manager" on page 189.

      View Switch Topology window

      View Endpoint Topology window

      View Logical Topology window

   c. Inspect the LEDs for the devices on the network and perform prescribed service procedures. Refer to "Interpreting LEDs" on page 122.

      Note: If this link has 4x switch ports configured to run at 12x with an octopus cable, for the LED behavior for groups of 4x switch ports connected to an octopus cable, refer to "Verifying static-12x mode connectivity" on page 67.

2. Look for installation problems:

   Perform this procedure after installation is complete.

   a. Open the IBM Network Manager on the HMC that is running the IBM Network Manager.
   b. Open each of the following windows within the IBM Network Manager, and verify that the connectivity is displayed according to the plan. If necessary, recable according to the plan.
      - View Switch Topology window
      - View Endpoint Topology window

- View Logical Topology window

3. Look for Configuration Issues:
   a. Open the IBM Network Manager on the HMC that is running the IBM Network Manager. Open each of the following windows within the IBM Network Manager, and verify that the connectivity is displayed according to the plan. If necessary, recable according to the plan.
      - View Logical Topology
      - View Switch Topology

   b. Verify that all adapters are configured and available to the operating systems. See "Verifying the installed InfiniBand network (fabric) in AIX or Linux" on page 67 for information about verifying that the adapters are configured and available.

      If an adapter is not configured and available to the operating system, use the operating system and the server documentation to configure the adapter. If the adapter cannot be configured, replace it.

   c. If this is a High Performance Cluster (HPC) network:
      1) Verify that all switches used for HPC applications are set to an LMC value of 2. Use the procedure in "Checking the logical identifier mask control" on page 147.
      2) Verify that the adapters are cabled in a balanced way so that the ports of each adapter are on separate subnets. Do this by checking the GID-prefixes in the Endpoint Topology View (choose **View Endpoint Topology** from the main IBM Network Manager window). If you find an incorrect prefix, do the following:
         a) If the adapter is cabled to the wrong switch, recable the adapter to the correct switch.
         b) Check the GID-prefixes for the switches using the procedure in "Checking GID-prefixes" on page 146.
         c) If the adapter is cabled correctly, and nothing is wrong with the GID-prefixes, it is possible that the network planning was not done properly, and more than one port on an adapter goes to the same subnet. If this is the case, and you have more than one subnet, the network planning must be redone, and the network must be recabled.

4. **This procedure ends here**

Suspected power problem

This procedure covers suspected switch power problems.

This procedure covers suspected switch power problems. For other suspected power problems, open Service Focal Point (SFP) on the controlling HMC for that device, then refer to the documentation for the associated server.

If you suspect a power problem with a switch, do the following:

1. Open SFP on the HMC on which the IBM Network Manager is running. Make sure that you perform the prescribed service actions for any switch power problem being reported in SFP on that HMC.
2. Go to the switch that has the suspected power problem and check the LEDs on the power modules. Perform any service procedure prescribed by "7048-270 or SFS7008P switch power supply LEDs" on page 126.
3. Verify that the power on the machine floor has been planned to meet the requirements of the cluster, and that no changes have been made to circuit layout since the time that the switch was operating without a power problem.
4. Refer to the operating power requirements in the appropriate switch manual for the operating requirements for the switch. To find the appropriate manual, refer to "InfiniBand switch reference information" on page 3.

5. If any inspection that you have made thus far leads you to believe that there is a problem with the power design on-site, contact your local planning specialists. Refer them to the appropriate manuals that define the operating environment requirements for the various devices in your cluster.
6. If you still suspect a power problem, consider that the problem may be caused by thermal issues. Refer to "Suspected thermal problem."

Suspected thermal problem

This procedure covers suspected switch thermal problems.

This procedure covers suspected switch thermal problems. For other suspected thermal problems, open Service Focal Point on the controlling HMC for that device, then refer to the documentation for the associated server.

If you suspect a thermal problem with a switch, do the following:

1. Go to Service Focal Point on the HMC on which the IBM Network Manager is running, and make sure that you perform the prescribed service actions for any switch thermal problem being reported in SFP on that HMC.
2. Depending on the switch type, perform one of the following actions:
   - For a 7048-120 or SFS7000P, the fans are in the same module as the power supplies. Go to the switch that has the suspected thermal problem and check the LEDs on the power modules. Perform any service procedure prescribed by "7048-270 or SFS7008P switch power supply LEDs" on page 126.
   - For a 7048-270 or SFS7008P, there are separate fan trays. Go to the switch that has the suspected thermal problem and check the LEDs on the fan trays, and perform any service procedure prescribed by "Switch fan tray LEDs" on page 127.
3. Inspect the site for airflow problems
   a. Look for any airflow obstructions that might be blocking floor tile vents. Look both above and below the floor.
   b. Ensure that the operating temperature is within temperature requirements. Refer to the appropriate switch manual, see "InfiniBand switch reference information" on page 3.
   c. If any inspection that you have made thus far leads you to believe that there is a problem with the airflow design on-site, contact your local planning specialists. Refer them to the appropriate manuals that define the operating environment requirements for the various devices in your cluster.
4. **This procedure ends here.**

Verifying adapters are configured and available

Provides links to procedure used when verifying adapter availability.

Depending on your operating system, select from one of the following links to verify that your GX 12X adapters are configured and available.

- AIX:
  - "Verifying the GX HCA connectivity in AIX" on page 67
- Linux:
  - "Verifying the GX HCA to InfiniBand fabric connectivity in Linux" on page 67

Checking the subnet manager

Procedures to check the subnet manager for the overall network and for individual switches.

Note: Each subnet must have a master subnet manager switch defined. If there is more than one switch in a subnet it is possible for switches to become isolated from other switches either because of cabling problems, or link failures. In such cases, corrective action must take place before the subnet manager configuration can be of use.

To check the entire network, check the overall subnet manager and then check the individual switch subnet managers:

1. To check the overall subnet manager, do the following:
   a. From the main IBM Network Manager window choose **View Management Properties**.
   b. Click the **Switch** tab.
   c. Verify the following for each switch in the list:
      - Note the location code and the switch name. Use **Rename Switch** to ensure that each switch has a meaningful name. It is suggested that the name reflects the switch's physical location, perhaps by using a frame number and slot within the frame as part of the name.
      - The switch is in the correct subnet by checking the GID-prefex, see "Checking GID-prefixes" on page 146. Corrective actions are described in "Switch is on the incorrect subnet" on page 144.
      - The switch has the correct LMC value; see "Checking the logical identifier mask control" on page 147.
      - The switch has the Subnet Manager Priority correctly set. Corrective actions are described in "Reordering Subnet Manager priority" on page 142.
      - The master status for the switch is appropriate, for example, Master, Standby, or blank. Corrective actions are described in "Reordering Subnet Manager priority" on page 142. Keep in mind that only one switch on a subnet should be the Master.
      - The Subnet Manager Version shows the correct code level. It is best if all subnet managers are at the same level of code. However, the subnet managers will operate correctly if their code levels are within one level of each other. Corrective actions are described in "Updating switch software" on page 144.
      - The Connectivity column indicates that the switch is responsive. If it is not, use the isolation procedure IBNSSLC; see "IBNSSLC" on page 116.
      - That the local times of all switches are reasonably close to one another. There is no function dependent on this time synchronization, but having them synchronized can be helpful in problem determination. Choose **Synchronize Time** and check the switch time relative to HMC time. If they are close enough, choose **No**. If you think they should be synchronized, choose **Yes**. Corrective actions are described in "Synchronizing Subnet Manager time with HMC time" on page 142.
2. Do the following to check the subnet manager information for individual switches:
   a. From the main IBM Network Manager window, open the **Switch Topology View** window.
   b. Find and select the switch on which you want to check the subnet manager. Make the selection at the chassis level.
   c. From the menu, click **Selected-Properties**.
   d. Click the **Subnet Manager** tab.
   e. Verify that the Subnet Manager location code matches the switch in which you believe it to be running. Only one switch in a subnet is configured to run the master Subnet Manager.
      - If this is the only switch in a subnet, the Subnet Manager location code matches its own location code.
      - If this switch is not the current master Subnet Manager, then the Subnet Manager location code will not be the same as this switch's location code, but rather it will match the location code for the master Subnet Manager's switch. Corrective action involves fixing broken links or cabling problems; see "Switch is on the incorrect subnet" on page 144.

Reordering Subnet Manager priority

Use this procedure to change the Subnet Manager order priority for a switch or a group of switches.

If you want to change the Subnet Manager priority for a switch or a group of switches, perform the following procedure:

1. From the main IBM Network Manager window, click **View Management Properties**.
2. Click the **Switch** tab.
3. For each switch that you want to change the Subnet Manager priority, perform the following procedure:
   a. Select the switch for which you want to change its Subnet Manager priority.
   b. Click the **Change SM Priority** button.
   c. In the **Change SM Priority** window:
      1) Click to **Enable SM** or **Disable SM** the Subnet Manager on the switch.
      2) If you chose to enable the Subnet Manager on the switch, fill in the **SM Priority** field with the correct Subnet Manager priority for this switch.
4. **This procedure ends here**.

Restarting the subnet manager

Use this procedure to restart the subnet manager.

To restart the subnet manager, perform the following procedure:

1. From the main IBM Network Manager GUI window, click **View Management Properties**.
2. Choose the **Switch** tab.
3. For each switch for which you want to restart the subnet manager, perform the following procedure:
   a. Select the switch for which you want to restart its Subnet Manager.
   b. Click the **Change SM Priority** button.
   c. In the Change SM Priority window click **Disable SM**.
   d. After the View Management Properties window's **Switch** tab reappears, select the same switch.
   e. Click the **Change SM Priority** button.
   f. In the Change SM Priority window, click **Enable SM** and enter the correct Subnet Manager priority level in the **SM Priority** field.
4. **This procedure ends here.**

Synchronizing Subnet Manager time with HMC time

Use this procedure when you want to synchronize the Subnet Manager time of a switch or group of switches.

To synchronize the Subnet Manager time for a switch, or group of switches, with the HMC time for the HMC that is running the IBM Network Manager, perform the following procedure:

1. From the main IBM Network Manager window, click **View Management Properties**.
2. Click the **Switch** tab.
3. For each switch for which you want to synchronize the Subnet Manager time with the HMC time, perform the following procedure:
   a. Choose the switch for which you want to synchronize the Subnet Manager time with the HMC time.
   b. Click the **Synchronize Time** button.

c. In the Synchronize Time window, verify that you selected the correct switch by checking the switch name and IP address.
d. Click the **Yes** or **No** button as appropriate.

4. **This procedure ends here.**

Location codes for machine types other than 7048

Use this procedure when servicing an InfiniBand switch that is a not a 7048 machine type.

Note: IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.

The recognized machine type switches have location codes for FRUs and chassis that begin with the following format:

- U7048.120.[chassis serial number]-P1 for a Topspin 120 Server Switch
- U7048.270.[chassis serial number]-P1 for a Topspin 270 Server Switch
- USFS7000Psssss.[chassis serial number]-P1 for a Cisco 7000 Server Switch
- USFS7000Psssss.[chassis serial number]-P1 for a Cisco 7008 Server Switch

The location codes may have more characters following the P1.

If the location code does not show a recognized machine type, consider the following possible reasons:

- A problem with the chassis VPD for the switch.
- The customer purchased a TS120 or a TS270 from IBM eServer hardware xSeries or IBM System x.
- The customer purchased a TS120 or a TS270 from a vendor.

If the location code is not a valid 7048 or SFS700*x* location code, perform the following procedure:

1. Find the switch in the cluster. The serial number for the switch precedes the P1 in the location code. For example, UTS12012X12XCPX.US1234567890-P1 indicates a serial number of US1234567890.
2. If there is a label on the switch that indicates that this is a 7048-120, 7048-270, SFS7000P, or SFS7008P, then the chassis VPD for this switch is incorrect, perform the following procedure. Otherwise, go to the next step.
   a. If this is a Topspin 120 Server Switch or Cisco 7000 Server Switch, replace the entire FRU chassis.

      Note: Use the RID tag procedures to document the original chassis serial number. This is required for service entitlement.

   b. If this is a Topspin 270 Server Switch or Cisco 7008 Server Switch switch, replace card number 17 which is the chassis ID module.
3. Determine if the customer has a service agreement with IBM to service this non-recognized machine type. If the customer has such an agreement, then continue service, and exit this procedure. Otherwise, go to the next step.
4. Inform the customer that the switch in question is not a recognized machine type.
5. If the customer has purchased the switch as a 7048 machine type, order a replacement switch that is correctly labeled as a 7048 machine type. Otherwise, go to the next step.
6. If the switch was purchased from IBM eServer hardware xSeries cluster 1350, ask the customer to contact IBM eServer hardware xSeries service.
7. If the switch was purchased from a vendor, ask the customer to contact the vendor's service provider.
8. **This procedure ends here.**

Checking switch software levels

Use this procedure to check on the switch software levels.

To check switch software levels, perform the following procedure:

1. From the main IBM Network Manager window, click **View Management Properties**.
2. Click the **Switch** tab.
3. Verify that the Subnet Manager version level is at the correct level.

Updating switch software

Use this procedure when updating the switch software.

If you want to update the software on a switch, or group of switches, perform the following procedure:

1. If you first need to check switch software levels, go to "Checking switch software levels."
2. From the main IBM Network Manager window, click **Update Switch Software**.
3. Select the switch(es) that you want to update.
4. Determine if the appropriate level of software is available on the HMC:
   a. Look at the Switch-Management Version panel for the desired of software.
   b. If it is available, proceed with step 5.
   c. If it is not available on this HMC, perform the following:
      1) Obtain a copy of the desired level of software from the switch manufacturer's Web site.
      2) Place the image in the /var/hsc/ibnm/TS/ directory.
      3) Alternatively, if you have a DVD with the correct level, you may insert the DVD into the drive for the HMC and click the **Import** button. Then, select the desired level available on the DVD and click **OK**.

      Note: If you believe that you have too many previous levels of switch software available on this HMC, use the Delete button to remove them.
5. Choose the desired level of software and click **OK**.

   The Update Switch Software Confirmation window opens.
6. Verify the list of switches for which you wish to update the software:
   a. If the list is correct, click **OK**. The Update Switch Software Install Summary window opens, proceed to step 7.
   b. If the list is incorrect, click **Cancel** and return to step 3.
7. If the Explanation column displays `Switch not responding`, perform the following:
   a. Verify that the power to the switch is powered on.
   b. Verify that the Ethernet connections and ports are operational from the HMC to the switch.
8. If any switch did not have the software installed correctly and you have attempted to fix the problem with the installation, return to step 3.

   If each switch indicates `OK` in the Explanation column, the installation is finished. Click the **OK** button to close the Update Switch Software Install Summary window.
9. **This procedure ends here.**

Switch is on the incorrect subnet

Use this procedure to determine why a switch is not on the correct subnet and has the incorrect GID-prefix.

If the switch does not have the expected GID-prefix, it is not going to behave as if it is on the correct subnet. There might be a subnet configuration problem or a cabling problem that caused this situation.

1. Verify that the subnet manager setup is correct. Ensure that the correct switches are setup as the master subnet manager. When you have only one switch per subnet, each switch should be the subnet manager for its own subnet.
2. Obtain the cable planning documentation.
3. Verify that the switch is cabled according to the planning documentation by:
   a. Checking the cabling to ensure that switches that should be isolated from one another are not connected. You can do this using the Switch Topology View and checking the connectivity as outlined in "Finding the other side of the link" on page 122. If you have only two switches and they are on separate subnets, a switch should never be attached to another switch. Otherwise, the cabling should match the original cable planning documentation. You are concerned only with switch to switch connectivity for this verification step.
   b. If you find a cabling problem, fix it, and recheck the GID-Prefix columns; see "Checking GID-prefixes" on page 146.
   c. If the switch is cabled correctly, verify that the planning documentation was generated properly. This task involves reviewing the planning steps again. If the cabling was not planned properly, rework the cabling according to a corrected plan.
4. **This procedure ends here.**

Understanding timestamp differences

Various sources of timestamp information to keep in mind when relating timestamp differences.

Because there are several time sources in a cluster, you might have to cross-reference the time found on a server and compare that to the time found on a switch. These sources might not be synchronized to one another or set to the same time-zone. Some examples of time sources are:

- A watch or a wall clock that you can reference when recording when you have performed an action.
- HMCs (there may be multiple HMCs in an cluster and they are not necessarily in sync)
- Service processors in a cluster of servers
- Switches
- A CSM management server

Whenever you are looking at a timestamp and trying to relate it to another timestamp, you might have to adjust for time-zone differences and time-base differences. A good practice when working in a clustered server environment is to check the current time from each source and note the differences.

Recovering from logical HCA configuration problems

Check the logical HCA configuration and correct the configuration using the IBM Network Manager's Logical Topology View window.

If your logical HCA configuration is not correct, use the Logical Topology View window in IBM Network Manager to modify the configuration.

Rebooting the entire switch chassis

Reboot the entire switch chassis when directed by problem isolation procedures.

Rebooting the entire switch chassis equates to rebooting the switch management software that runs in the switch chassis. Perform the following procedure:

Note: If you need to reboot a switch chassis, all resources within it will be temporarily unavailable to perform their normal functions. This might result in errors that are reported to Service Focal Point.

1. From the main IBM Network Manager window, select **View Management Properties**.
2. Select the **Switch tab**.
3. Select the switch that you want to reboot.
4. Select **Reset Switch Box**.

Adjusting Firewall Parameters for SNMP Traps

In order for IBM NM to receive critical SNMP traps from the switches, the SNMP port must be opened on the HMC firewall.

From the HMC GUI, in the navigation area:

1. Expand the HMC that controls the switch from which you wish to receive error events.
2. Expand the **HMC Management** section.
3. Select **HMC Configuration** and the menu will appear in the right frame.
4. Select **Customize Network Settings**
5. Click on the **LAN Adapters** tab.
6. Select the LAN adapter that is on the same network as your switch.
7. Click on **Details**....
8. Click on the **Firewall** tab.
9. In the upper box, select the application that is named SNMP Traps, then click on **Allow Incoming**.
10. Verify that SNMP Traps has been added to the list of allowed applications in the lower box.
11. Click **OK**.
12. Click **OK**.
13. You must now reboot your HMC in order for the changes to be applied.

This ends the procedure.

GID-prefix procedures

GID-prefixes are used to identify groups of devices. IBM System p5 or eServer p5 uses GID-Prefixes to identify subnets.

Checking GID-prefixes

Check the GID-prefixes used to identify subnets.

To check GID-prefixes, use the following procedure:

1. From the main IBM Network Manager GUI window choose **View Management Properties**.
2. Click the **Switch** tab.
3. Check the **Assigned GID-Prefix** column:
   - If the Assigned GID-Prefix is correct, go to the next step.
   - If the Assigned GID-Prefix is incorrect, use the following procedure to set it to the correct value: “Setting GID prefixes” on page 147
4. Check the **SMM-GID-Prefix** column:
   - If the **SMM-GID-Prefix** column is correct, **This procedure ends here**.
   - If the **SMM-GID-Prefix** column is incorrect, see “Switch is on the incorrect subnet” on page 144. **This procedure ends here.**

Setting GID prefixes

GID-prefixes are used to identify subnets.

To set the GID-Prefix use the following procedure:

1. From the main IBM Network Manager GUI window choose **View Management Properties**.
2. Click the **Switch** tab.
3. Select the switch for which you wish to change the GID-prefix.
4. Click **Set Subnet Parameters**.
5. Under **Assign GID-Prefix**, check under **Assign from Available GID-Prefixes** to determine whether the prefix that you want to assign already exists in the cluster.
6. Does the GID-prefix already exist?

   Yes Select the desired GID-prefix from Assign from Available GID-Prefixes list.

   No Enter the desired GID-prefix in the Assign New GID-Prefix field.

   Important:
   Before clicking **OK**, check the Assign LMC value.

   If it is correct, click **OK** . **This procedure ends here.**

   If it is incorrect, then use the procedure in "Setting the location identifier mask control" on page 148.
7. Reboot all servers on the subnet to recognize the new GID-prefix. If you are changing multiple GID-prefixes, change them all before rebooting servers.

Logical identifier mask control procedures

The LID (logical identifer) mask control (LMC) effects the number of possible LIDs for a point in a subnet.

Each point in a network must have at least one logical ID (LID). The quantity of LIDs for a point (which is used for route addressing) indicates the quantity of routes that can be used to access the point in the network. The default value for the LMC in a subnet is to have one route per point. In the default case, each point on the subnet would have one LID.

For IBM System p5 and eServer p5 High Performance Computing (HPC), servers require four (4) routes per point. In this case there would be 4 LIDs per point on the network.

The LMC value is equal to the exponent for a power of two that will result in the desired number of LIDs. To calculate the default value for the LMC of one (for 1 route per point on the network) raise 2 to the power of 0 (2^0), which results in 1. For HPC, because four routes are required, raise 2 to the power of 2 (2^2), which results in 4.

In a cluster with only one switch per subnet, that switch must contain the Subnet Manager Master (SMM). For each subnet, there is an assigned LMC value and a SMM-LMC value. The subnet manager has control of the actual LMC value used by all switches in a subnet. The subnet manager uses its switch's Assigned LMC value to set the LMC value that will be used by the subnet. Therefore, the assigned LMC and SMM-LMC should always be the same, unless there is a cabling problem that results in switches in different subnets being cabled together.

Checking the logical identifier mask control

A procedure for checking the LID (logical identifier) mask control (LMC)

To check GID-Prefixes, use the following procedure:

1. From the main IBM Network Manager GUI window choose **View Management Properties**.
2. Click the **Switch** tab.
3. Check the Assigned LMC column:
   - If the assigned LMC is correct, go to the next step.
   - If the assigned LMC is incorrect, use the following procedure to set it to the correct value: see "Setting the location identifier mask control."
4. Check the SMM-LMC column:
   - If the SMM-LMC column is correct, **this procedure ends here.**
   - If the SMM-LMC column is incorrect, use "Switch is on the incorrect subnet" on page 144.

Setting the location identifier mask control

Set the LID (location identifier) mask control (LMC) using the IBM Network Manager GUI window.

To set the LMC do the following:

1. From the main IBM Network Manager GUI window click **View Management Properties**.
2. Click the **Switch** tab.
3. Select the switch for which you wish to change the LMC.
4. Choose **Set Subnet Parameters**.
5. Under **Assign LMC**, select the desired LMC value. Recall that for HPC applications the LMC should be set to two (2). All other applications can have the default LMC value of zero (0).
6. Before clicking **OK**, verify the GID-prefix values and make sure they are assigned properly. See "Setting GID prefixes" on page 147. Because you are already viewing the Set Subnet Parameters window, you can skip forward to the appropriate steps in the procedure.
7. After changing the LMC, you must reboot all servers on the subnet to recognize the new LMC. If you are changing multiple LMCs, change them all before rebooting servers. **This procedure ends here.**

Repairs when using IBM Network Manager

Procedures used to minimize extraneous events that can be caused by repair actions and also how to verify repairs using IBM Network Manager.

Note: IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.

Perform the "Repair preparation" procedure before beginning FRU replacement.

If you disabled a link to prepare for "Repair preparation," you must re-enable it as in the "Re-enable link(s) after repair" on page 149 procedure.

Perform the "Repair verification" on page 150 procedure after the FRU has been replaced and brought back into service.

Repair preparation

If you are making a repair that will cause a link to go down, it is strongly suggested that you disable the link from the switch port before performing service. This should not have to be done if you are replacing a redundant FRU such as a power supply, fan tray, core fabric controller module, or I/O Management module that has an operating alternative FRU.

1. Record the location of the FRU that is to be repaired.

2. If you are replacing an HCA, find the HCA in the Endpoint Topology view (choose **View Endpoint Topology** from the IBM Network Manager main window), and determine the switch neighbors.
3. From the IBM Network Manager main window choose **View Switch Topology**.
4. Find the switch which contains the port. If you are replacing a 7048-120 or SFS7000P, or a LIM in a 7048-270 or SFS7008P go to step 5.

   If you are replacing a fabric controller in a 7048-270 or SFS7008P, go to step 6.

   Note:

   a. If you must disable ports on separate switches, remember that switches that are on different subnets have different GID-prefixes and you must choose the appropriate GID-prefix to find the switch.

   b. If a switch port has a 4x connector from an octopus cable attached to it, carefully record to which port each of the 4x connectors of the octopus cable connects. You must replace them in the correct order. If you do not do this, you will need to determine the configuration again using the procedure in "Planning for InfiniBand networks" on page 5. For more on behavior of 12x connections see "Planning for InfiniBand networks" on page 5.

   a. Expand the card that contains the port.
   b. Select the port.
   c. From the menu, choose **Selected-Administrate-Disable**.
   d. **This procedure ends here.**
5. If you are replacing a 7048-120 or SFS7000P, or a LIM in a 7048-270 or SFS7008P, select the card you are replacing and choose **Selected-Administrate-Disable**. **This procedure ends here.**
6. If you are replacing a fabric controller in a 7048-270 or SFS7008P, use the following to determine which LIM cards should be disabled to avoid extraneous serviceable events."
   a. Determine which LIMs, if any, are associated with the fabric controller using "Plugging Rules for 7048-270 or SFS7008P switch line interface modules and fabric controllers" on page 152.
   b. Select each associated LIM and choose **Selected-Administrate-Disable**.
   c. **This procedure ends here.**

Re-enable link(s) after repair

If you disabled a link to prepare for the repair procedure, you might have to re-enable it to bring it back online. You must do this before the LEDs will go to a good state.

Note: If you replaced a 7048-120 or SFS7000P switch, then you replaced the major components of the switch so there is no memory of links being disabled.

Perform the following procedure:

1. Record the location of the FRU that is to be repaired.
2. If you are replacing an HCA, find the HCA in the Endpoint Topology view (choose **View Endpoint Topology** from the IBM Network Manager main window), and determine the switch neighbors.
   - Checking the cable planning documentation
   - Checking the cable labels
   - Tracing the cable under the floor.
3. If you replaced a switch card or chassis that had octopus cables connected to it, you will need to configure all of the ports on the new switch card(s) that are acting as a group in a 12x link using an octopus cable. See the procedure in "Configuring Static-12x groups" on page 20. For more on behavior of 12x connections see "Planning for InfiniBand networks" on page 5.
4. From the IBM Network Manager main window choose **View Switch Topology**.
5. Find the switch that contains the port. If you are replacing a 7048-120 or SFS7000P, or a LIM in a 7048-270 or SFS7008P, go to step 6 on page 150.

If you are replacing a fabric controller in a 7048-270 or SFS7008P, go to step 7.

Note: If you must enable ports on separate switches, remember that switches that are on different subnets have different GID-prefixes and you must choose the appropriate GID-prefix to find the switch.

a. Expand the card that contains the port.
b. Select the port.
c. From the menu, choose **Selected-Administer-Enable**.

Note: It is possible that the card is already enabled, and you will not be able to select **Selected-Administer-Enable**. If this is true, then there is no need to perform this step.

d. **This procedure ends here.**

6. If you are replacing a 7048-120 or SFS7000P, or a LIM in a 7048-270 or SFS7008P, select the card you are replacing and choose **Selected-Administer-Enable**.

Note: It is possible that the card is already enabled, and you will not be able to select **Selected-Administer-Enable**. If this is true, then there is no need to perform this step. **This procedure ends here.**

7. If you are replacing a fabric controller in a 7048-270 or SFS7008P, use the following to determine which LIM cards should be disabled to avoid extraneous serviceable events.
   a. Determine which LIMs, if any, are associated with the fabric controller using "Plugging Rules for 7048-270 or SFS7008P switch line interface modules and fabric controllers" on page 152.
   b. Select each associated LIM and choose **Selected-Administer-Enable**.

Repair verification

When repairing a failing FRU in the switch network, it is important to verify that the repair has restored network service with that FRU. Perform the following procedure for repair verification:

1. Ensure that if you disabled a link in preparation for repair that you re-enable it before you proceed with the remainder of this procedure.
2. Perform any repair verification procedures recommended by the repair procedures for a specific FRU.

Note:

a. Because repair actions can be disruptive to the network not only in the direct replacement of a FRU, but also in an indirect manner through inadvertently unseating cables and other unintended consequences, the remainder of this procedure has you check connectivity to switches and the service subsystem.
b. Always check the LEDs for proper status when you have replaced a FRU. The LEDs are your first indication if there is a problem. See "Interpreting LEDs" on page 122.
c. If a switch is unresponsive during a repair procedure, serviceable events may be lost during the time it is unresponsive. Check the port status of all ports on the switch after a repair action.

3. If the FRU was an HCA and you did not disturb the connected switch, go to step 4. Otherwise, do the following procedure to check the switch connectivity to the service network:
   a. From the IBM Network Manager main GUI window, choose **View Switch Topology**.
   b. Verify that the switch Connectivity status is Responsive. If it is not Responsive, perform the procedure in "IBNSSLC" on page 116.
4. Check the switch connectivity to the InfiniBand network:
   a. In the Switch Topology view, find the switch chassis icon that contains the FRU and click on the icon to expand it.
   b. Expand all the card and port levels beneath the switch chassis.

c. Verify that all ports have the correct neighbor location according to the cabling plan. If you find a misplaced cable, re-cable according to the cabling plan.
d. Verify that all ports which are connected to another device indicate that they are Active. If any port that is expected to be Active is not Active, perform the procedure in "Port status in the switch topology window" on page 201.

Note: If the switch port has a 4x connector from an octopus cable attached to it, verify connectivity using the procedure found in "Verifying static-12x mode connectivity" on page 67.

5. If the FRU was a power card, perform the following make sure that you have verified that the FRU is operational by checking the LEDs as indicated in "7048-270 or SFS7008P switch power supply LEDs" on page 126, then do the following:
   a. Assuming you are still in the Switch Topology View, find the switch chassis that contains the FRU and click on it.
   b. From the menu, choose **Selected-Environmentals**.
   c. Verify that the FRU is listed in the Power Supply panel and that it has good status as indicated in "Switch Environmental Status" on page 205.

   Note: Power supplies have a location code that ends in -E1 or -E2.

   If this is a 7048-120 or SFS7000P, also verify that the fans are operating properly by checking the Fans panels as indicated in "Switch Environmental Status" on page 205.

   Note: Fan-id 1 and 2 are on supply E1 and fan-id 3 and 4 are on supply E2.
6. If the FRU was a fan tray, make sure the FRU is operational by checking the LEDs, as indicated in "Switch fan tray LEDs" on page 127, then perform the following:
   a. From the IBM Network Manager main GUI window, choose **View Switch Topology**.
   b. Find the switch chassis that contains the FRU and click on it.
   c. From the menu, choose **Selected-Environmentals**.
   d. Verify that the FRU is listed in the Fans panel and that it has good status as indicated in "Switch Environmental Status" on page 205.

   Note: Fan trays have a location code that ends in -A1 or -A2. Also, there are two fans on each fan tray. Fan-id 1 and 2 are on -A1, and Fan-id 3 and 4 are on -A2.

Perform the following procedure after replacing an HCA:

1. Verify that the server is on the service network:
   a. From the IBM Network Manager main window, choose **View Endpoint Topology**.
   b. Verify that the server that contains the HCA FRU is visible in the IBM Network Manager. If it is not, perform the following:
      1) If no servers are visible, check the Ethernet connection from the service network into the HMC running IBM Network Manager.
      2) Check the Ethernet connection into the frame and the server's service processor.
      3) Check the Ethernet network that supports the service network.
      4) Check the wall power to the server or to the frame that contains the server.
2. Verify that the server that contains the HCA FRU is powered on by checking the Power column. If it is not, power it on.
3. Verify the HCA connectivity to the InfiniBand network:
   a. Assuming that you are in the Endpoint Topology view, expand the server that contains the HCA.
   b. Expand the **HCA**.

c. Check the status for each port on the HCA as indicated in "Adapter status in End-Point Topology window" on page 190. Each port that has a cable should indicate it is active.

4. Verify the logical topology information for the HCA:
   a. Make sure that the LPARs associated with the repaired HCA are booted.
   b. From the IBM Network Manager main window, choose **View Logical Topology**.
   c. Go to the server that contains the repaired HCA.
   d. Expand the server.
   e. Expand the physical host channel adapter, **PHCA**.
   f. Expand the logical host channel adapter, **LHCA**.
   g. Check the status for the LHCAs and logical switches (LSWs) as indicated in "Logical topology window status" on page 199.

      Note: All LHCA ports should indicate they are active. Each LSW port 1 status corresponds to a port to the switch network. Only the LSW ports that have a cable connected to them will indicate they are active.

Plugging Rules for 7048-270 or SFS7008P switch line interface modules and fabric controllers

The relation of line interface modules (LIMs) and fabric controllers with one another on a 7048-270 or SFS7008P.

Because of the way the backplane is wired between LIMs and the fabric controllers in a switch, certain fabric controllers support certain LIMs. Also, LIMs are populated in pairs and they must be the same LIM type. If you cannot get connectivity to a switch to work, or LIMs appear to be missing, check that the following rules have been met.

LIM to fabric controller groupings are as follows:

- LIM slots 1 and 2 are connected to fabric controller slot 9
- LIM slots 3 and 4 are connected to fabric controller slot 10
- LIM slots 5 and 6 are connected to fabric controller slot 13
- LIM slots 7 and 8 are connected to fabric controller slot 14

LIM pairs must be the same type. LIM pairings are as follows:

- slots 1 and 2 must be the same type
- slots 3 and 4 must be the same type
- slots 5 and 6 must be the same type
- slots 7 and 8 must be the same type

Disappearing IBM Network Manager windows

Disappearing IBM Network Manager windows may be caused by a restart of the IBM Network Manager daemon.

Perform the following procedure:

1. Reopen the window after several minutes.
2. You might see indications of missing devices or unknown status because the IBM Network Manager is attempting to find network devices that have a problem. Wait several minutes for IBM Network Manager to update correctly.
3. If after 10 minutes have passed, updates have not occurred, restart the IBM Network Manager.

4. If there are still missing devices or status that is unknown, perform the appropriate status procedures. If you still experience a problem, make sure there are no hardware issues and then call IBM software support.

Verifying Static12x or 4x configuration to a port

Verify static 12x or 4x configuration to a port.

To verify how a port is configured for 4x or 12x operation, perform the following:

1. Open the IBM Network Manager Overview window.
2. Select **View Switch Topology**.
3. Select the appropriate GID-prefix to display the switch that contains the port.
4. Expand the switch and card that contains the port.
5. Select the port.
6. Choose **Selected-Properties**.
7. Choose the System tab.
8. Note the following properties: Width supported, width enabled, width active.

   Width supported
   is the list of supported speeds that this can support alone and outside of a group of ports.

   Wdith enabled
   is the chosen speed for this port. If it is 12x, and this is the lowest numbered port in a group of three ports as described in "Planning for InfiniBand networks" on page 5, then the group is expected to be used to connect to a single 12x device using an octopus cable.

   Width active
   is the actual speed for this port. It is expected to normally be 4x or 12x. If it is 1x, then the port has been configured to auto-negotiate. If it is 12x, and this is the lowest numbered port in a group of three ports as described in "Planning for InfiniBand networks" on page 5, then the group is expected to be used to connect to a single 12x device using an octopus cable.

Determining faulty fabric controller cards versus faulty LIM cards

A faulty fabric controller card can cause multiple switch ports to fail on a 7048-270 or SFS7008P.

If you have multiple ports failing, use the "Plugging Rules for 7048-270 or SFS7008P switch line interface modules and fabric controllers" on page 152 to determine if the ports are associated with a LIM pair. Then, use the following procedure:

1. If the ports are all on the same LIM, replace the LIM card. If this does not fix the problem, replace the associated fabric controller card as determined in "Plugging Rules for 7048-270 or SFS7008P switch line interface modules and fabric controllers" on page 152.
2. If the ports are across two LIM cards in the same LIM pair, replace the associated fabric controller card as determined in "Plugging Rules for 7048-270 or SFS7008P switch line interface modules and fabric controllers" on page 152.

Missing LIM(s) or LIM ports

Determine problems with missing LIM(s) or LIM ports

There are several reasons why a LIM may not have all of its ports listed in the Switch Topology View. They are:

- An unpopulated LIM slot
- A poorly seated LIM

- A faulty LIM
- A faulty fabric controller card

If you have one or more missing LIM or LIM ports that should be populated perform the following procedure:

1. As you take an action on a FRU by either reseating or replacing it, remember to check the Switch Topology View to see if the action has resolved the problem.
2. Verify that all LIMs display all of their ports when they are expanded in the Switch Topology View. A 4x switch has 12 ports per LIM.
3. If you are missing all of the LIMs in a switch or an entire switch, do the following checks:
   - Check the chassis LEDs for the switch
   - Check the switch's Ethernet connection to the service subsystem.
   - Check the other service network's connections.
   - Check the Ethernet routers and switches for the service subsystem.
4. Make sure that the LIMs adhere to the plugging rules in "Plugging Rules for 7048-270 or SFS7008P switch line interface modules and fabric controllers" on page 152.
5. If more than one LIM is missing its ports or is not visible in the Switch Topology View, record all of the LIMs that are missing their ports or are not visible, and indicate if they are members of the same LIM pair. Also, record the associated fabric controller(s) for the LIM(s). The LIM pairs are listed below:
   - LIMs 1 and 2, which are controlled by the fabric controller in slot 9.
   - LIMs 3 and 4, which are controlled by the fabric controller in slot 10.
   - LIMs 5 and 6, which are controlled by the fabric controller in slot 13.
   - LIMs 7 and 8, which are controlled by the fabric controller in slot 14.

   Note: Slot mappings are found in "Switch FRU Locations for the 7048-270 or SFS7008P" on page 220.
6. Verify that the LIM slot is populated and that the LEDs for all LIMs and fabric controllers are indicating no issue. Interpretation of the LEDs and proper recovery procedures are found, below.
   - Interpretation of LIM Status LEDs and the proper procedures for bad status are found in "Line Interface Module status indicator" on page 128.
   - Interpretation of Fabric Controller LEDs and the proper procedures for bad status are found in "Fabric controller LEDs" on page 130.
7. Verify that the LIM is properly seated. It is hot-pluggable, so you may reseat it and check the Switch Topology View again. Be careful not to dislodge cables on any other LIMs that are operational.

   Note: If more than one LIM is missing its ports and they are members of a LIM pair (LIMs 1 and 2, LIMs 3 and 4, LIMs 5 and 6, or LIMs 7 and 8), then it is likely that there is an issue with a fabric controller. You may choose to perform the above step, or you may proceed to the next step and make a note that you did perform the step and return to it if the remaining steps of this procedure do not resolve the problem.
8. If only one LIM in a LIM pair is missing ports, perform the following procedure; otherwise proceed to the next step.
9. At this point, a fabric controller is the likely problem. Perform the following procedure on the fabric controller that you recorded in step 5 as being associated with the problem LIM(s).
   - Reseat the fabric controller and check the Switch Topology View to see if the problem has been fixed.
   - If the problem still exists, replace the fabric controller and check the Switch Topology View to see if the problem has been fixed.
10. Verify that you have checked all of the following potential causes for missing LIM ports:
    - The LIM(s) are populated

- The LIM and fabric controller LEDs state
- The LIM(s) are properly seated
- The fabric controllers are properly seated
- The LIM(s) are not faulty
- The fabric controller(s) are not faulty
- You have verified that the problem still is visible in the Switch Topology View.

11. If you have verified all of the above potential causes for missing LIM ports, it is likely that the backplane is the problem. Contact your next level of support and request a replacement chassis which contains the backplane.

Administrative procedures for InfiniBand switches

Procedures in this section are listed in a table that links to various administrative procedures for switches.

Note: IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.

At this time, the IBM Network Manager does not provide access to all of the administrative functions for InfiniBand switches. Some administrative functions for a switch are accessed through the serial administrative port that is located on the switch.

Note: Procedures that use administrative port access for a switch will be disruptive to the network.

Description	Switch administrative procedure
Access switch administrative functions to connect to a serial port on the switch chassis.	"Accessing a serial port on a switch"
Access the serial port on a 7048-120 or a SFS7000P switch.	"Accessing a 7048-120 or a SFS7000P serial port" on page 158
Access the serial port on a 7048-270 or a SFS7008P switch.	"Accessing a 7048-270 or a SFS7008P serial port" on page 158
Access the InfiniBand switch command line interface (CLI).	"Accessing an InfiniBand switch command line interface" on page 158
Cross-reference to "Accessing a serial port on a switch"	"Reviewing the ts_log file" on page 159
Setting IP addressing in switches	"Setting IP addressing in switches" on page 159
Shutting down the switch card function and then starting it up again.	"Rebooting switch cards" on page 160
Run diagnostics on cards using the InfiniBand switch command line interface.	Running Switch Card Diagnostics. Refer to "InfiniBand switch card diagnostics" on page 173
The IBM Network Manager works with database synchronization enabled. If this was disabled, use this procedure to enable database synchronization.	"Setting database synchronization" on page 161
Use this procedure to adjust one of the timeout values for database synchronization.	"Adjusting database-synchronization timeout" on page 161
Use this procedure to access the FRU Identification LEDs on a Topspin switch.	"Accessing FRU Identification LEDs" on page 162
To initiate bug reports.	"Filing bug reports" on page 163

Accessing a serial port on a switch

Use this procedure to access administrative functions for the switch.

To access switch administrative functions, a connection to the serial port on the switch chassis is required. The procedure for connecting to a serial port depends on the model. Select one of the following serial-port-connection procedures.

- "Accessing a 7048-120 or a SFS7000P serial port" on page 158

© Copyright IBM Corp. 2004, 2007

- "Accessing a 7048-270 or a SFS7008P serial port"

Accessing a 7048-120 or a SFS7000P serial port

Use this procedure to access the serial port on a 7048-120 or SFS7000P switch.

1. Connect the cable from the 7048-120 or SFS7000P serial port to your personal computer; use the straight-through (M or F) serial cable, which is provided in the switch package. For detailed information about how to connect the serial cable, see the documentation included with the serial cable kit.
2. Open a terminal emulation window using a program such as HyperTerminal for Windows. Set your terminal parameters to the following:
   - Baud: 9600 bps
   - Data Bits: 8
   - Parity: None
   - Stop Bits: 1
   - Flow control: None
3. **This procedure ends here.**

Accessing a 7048-270 or a SFS7008P serial port

Use the following procedure to access the serial port on a 7048-270 or a SFS7008P switch.

1. Attach the RJ-45 console cables from the cable kit that is provided.
   a. Connect the cables to the InfiniBand switch chassis serial-console port on the management-interface card.

      Note: If you have two management interface cards, ensure that you connect to the left management interface module (slot 15), which is the primary management card upon initial boot. The serial console port is labeled **10101**.

   b. Connect the other end of the serial cable to your terminal server or management workstation.

      Note: For detailed information about how to connect the serial-console cable, see the documentation included with the serial-cable kit.
2. Open a terminal emulation window using a program such as HyperTerminal for Windows. Set your terminal parameters to the following:
   - Baud: 9600 bps
   - Data Bits: 8
   - Parity: None
   - Stop Bits: 1
   - Flow control: None
3. **This procedure ends here.**

Accessing an InfiniBand switch command line interface

Use this procedure access the InfiniBand switch Command Line Interface (CLI).

The command line interface (CLI) is documented in detail in the Topspin CLI document; refer to the Topspin documentation found in "InfiniBand switch reference information" on page 3.

In order to access the Topspin Switch CLI, you must first access the Topspin Switch administrative port using the procedure described in "Accessing a serial port on a switch" on page 157.

1. After you are connected to the serial port of the switch, the CLI login prompt is displayed.

Note:

- For service purposes, use administrative (admin) authority. Obtain the admin login password from the system administrator. If local security guidelines prohibit this, discuss with the system administrator which commands must be run and determine if they are available under a guest login. If they are not, you must get the system administrator's help to run the commands. Most procedures require admin login privileges.
- If you require Privileged Execute mode for any commands, then the guest login will not be sufficient. If one of the first commands that you require is the **enable** command, you must have the admin login capability.

2. Log in using the provided login ID and password.
3. Run the required commands from the CLI login prompt
4. To exit the CLI session, type `exit`.
5. **This procedure ends here.**

Reviewing the ts_log file

If you are instructed to collect or review the ts_log file, use this procedure.

Note: FTP must be enabled on the HMC.

To collect the ts_log file, run the following command: /opt/hsc/bin/tslog.snap –hmc_ip *ip address of the HMC* –switch_ip [*ip address of the switch*] – out_file_name *filename to store locally on HMC*

The output file name includes PMH and branch office numbers and a timestamp: `tslog.`*`PMH-branch.yearmonthday-hourminute`*

Setting IP addressing in switches

Setting up IP addressing for switch networks with single or multiple HMCs.

To setup the correct IP addressing for switches, if you have access to the service network and can **telnet** directly into the switches, it is still advisable to set IP addressing parameters using the switch's Command Line Interface (CLI).

Do one of the following for your configuration:

- If you have only a single HMC, or an HMC with a redundant partner managing the same devices, then you must set up the switch for DHCP service for its IP addressing. Go to "Setting IP addressing for a switch network with a single HMC."
- If you have multiple HMCs that manage different devices, then you should be using a Cluster-Ready Hardware Server and must setup the switch with a fixed IP address. Go to "Setting IP addressing for a network with multiple HMCs" on page 160.

Setting IP addressing for a switch network with a single HMC

To setup a switch for DHCP service for IP addressing, do the following:

1. Access the administrative port for the switch using the procedure described in "Accessing a serial port on a switch" on page 157.
2. Log into the CLI using the procedure described in "Accessing an InfiniBand switch command line interface" on page 158. Admin authority is required.
3. Enter `enable`
4. Enter `config`
5. Enter `interface mgmt-ethernet`
6. Enter `addr-option dhcp`

7. Enter `exit`
8. Enter `exit`
9. To save this change for subsequent boots of the switch, enter: `copy running-config startup-config`

   Note: Do not logoff the CLI before completing this step. If you logoff, you will have to repeat this procedure anytime the switch is booted.
10. Enter `enable`.
11. Enter `reload no-failover`
12. Enter `YES` when asked to `Proceed to reload?`. The switch will now reboot. For clusters with GX HCAs, the HMC assigns an IP address to the switch.
13. Terminate your terminal session and disconnect the personal computer.

Setting IP addressing for a network with multiple HMCs

A fixed IP address is required for switches in a network with multiple HMCs. To set up a switch to have a fixed IP address, do the following:

1. Access the administrative port for each switch using the procedure described in "Accessing a serial port on a switch" on page 157.
2. Log into the CLI using the procedure described in "Accessing an InfiniBand switch command line interface" on page 158. You will require admin authority.
3. Enter `enable`
4. Enter `config`
5. Enter `interface mgmt-ethernet`
6. Enter `addr-option static`
7. Enter `ip [ip-address]`
8. Enter `exit`.
9. Enter `exit`.
10. Enter `exit`.

Rebooting switch cards

Problem isolation procedures can use rebooting of switch cards to help diagnose a problem.

In certain isolation procedures, you might be requested to reboot a switch card. The process requires shutting down the card function and then restarting the card. Perform the following procedure:

Note: If you need to boot a switch card, the card will be temporarily unavailable to perform its normal functions. This might result in downstream errors that are reported to Service Focal Point. Therefore, be prepared to close out Service Focal Point events reported by devices that are connected to the card, and by devices that are monitoring the card.

1. Access the administrative port for the specific switch by using the procedure described in "Accessing a serial port on a switch" on page 157.
2. With admin authority, log in to the CLI using the procedure described in "Accessing an InfiniBand switch command line interface" on page 158.
3. Use the following procedure for disabling and enabling switch cards. For details, see the Topspin CLI documentation in "InfiniBand switch reference information" on page 3.
   a. In User Exec mode, enter the **enable** command to enter Privileged Exec mode.
   b. Enter the **configure terminal** command to enter Global Configuration mode
   c. Enter the **card** ***card-selection*** command and specify the card or cards that you want to enable. The *card-selection* can be a comma-delimited list of cards, for example 1, 2, 5, 10.
   d. If the switch model is a 120 or 700, enter `reload`.

If the switch model is a 270 or 708, enter `action reset`.

Note: If the **action reset** command returns an error and indicates that you are to use the **reload** command, you may do this, but *all* cards in the chassis will be reset. Therefore, consider deferring this action to a time that will minimize impact to the customer's network.

4. Wait for a few minutes.
5. Check the card LEDs to verify that the switch reset without errors; see "Interpreting LEDs" on page 122.

   Note: If you used the **reload** command in the preceding step, check the LEDs on all of the cards in the chassis.
6. To verify that you have good status, check the IBM Network Manager Switch Topology View window.
7. Check Service Focal Point to see if any serviceable events were logged.
8. **This procedure ends here.**

Setting database synchronization

The database synchronization value must be correctly set depending upon your switch configuration.

To check the current value setting, and to reset the value if necessary, perform the following procedure:

1. To find the subnet manager master switch, use the procedure found in "Checking the subnet manager" on page 140.
2. Access the administrative port for the subnet manager master switch in question using the procedure described in "Accessing a serial port on a switch" on page 157.
3. With Admin authority, log in to the Command Line Interface (CLI) using the procedure described in "Accessing an InfiniBand switch command line interface" on page 158.
4. From the command line, query the current value for database synchronization by entering: `show ib sm db-sync`
5. The returned value needs to match your existing switch configuration:
   a. If you have a single switch per subnet the value returned should be `FALSE`. If the value returned is *not* `FALSE`, issue the following command by entering: `no ib sm db-sync subnet-prefix prefix enable`
   b. If you have multiple switches per subnet then the value returned should be `TRUE`. If the value returned is *not* `TRUE`, issue the following command by entering: `ib sm db-sync subnet-prefix prefix enable`
6. If you had to set the value, check that the command was successful by reentering: `show ib sm db-sync`
7. **This procedure ends here.**

Adjusting database-synchronization timeout

Adjust one of the timeout values for database synchronization.

To adjust one of the timeout values for database synchronization, do the following:

1. Access the administrative port for the subnet-manager master switch using the procedure described in "Accessing a serial port on a switch" on page 157.
2. With admin authority, log in to the Command Line Interface (CLI). See "Accessing an InfiniBand switch command line interface" on page 158.
3. On the subnet-manager master switch, query the desired value (session-timeout or cold-sync-timeout) Use the **show ib sm db-sync** command as described in the Topspin CLI documentation. See "InfiniBand switch reference information" on page 3.

If you need to find the subnet-manager master switch use the procedure found in "Checking the subnet manager" on page 140.

4. Use the **ib sm db-sync** command as described in the Topspin CLI documentation. See "InfiniBand switch reference information" on page 3. Set the session-timeout value and the cold-sync-timeout value to a higher value than the values found in the preceding query.
5. **This procedure ends here.**

Accessing FRU Identification LEDs

Use this procedure to access the FRU Identification LEDs on a Topspin switch.

If the IBM Network Manager does not provide access to a FRU identification LED, you will have to connect a terminal to the Topspin switch serial port and use the command line interface (CLI) to perform the desired actions. To access the FRU Identification LEDs on a Topspin switch, use the following procedure.

Note:

- To turn an identification LED on or off, use the IBM Network Manager. This helps you to avoid overriding the current status information for the device LED.
- To turn off an identification LED, and IBM Network Manager does not provide the function, you can use the following procedure. Instead of checking to see if the LED is on after entering the `start` command, you can immediately follow `start` with the `stop` command.

1. Access the administrative port for the subnet manager master switch using the procedure described in "Accessing a serial port on a switch" on page 157.
2. With admin authority, log in to the command line interface using the procedure described in "Accessing an InfiniBand switch command line interface" on page 158.
   a. In User Exec mode, enter the `enable` command to enter Privileged Exec mode.
   b. Enter the `configure terminal` command to enter Global Configuration mode.
3. In the switch command line interface, you will use the `diagnostic` command to enter diagnostic mode. For details on the `diagnostic` command, see the Topspin CLI documentation on the "InfiniBand switch reference information" on page 3. Use one of the following options to choose the appropriate card:
   a. For LIMs, or Fabric Control Module use: `diagnostic card` *slot_number*.
   b. For power supplies (or combined power/fan modules), use: `diagnostics power-supply` *supply*.
   c. For fans, use: `diagnostic fan` *fan-number*.
   d. For the chassis module, use: `diagnostic chassis`.
   e. For the Ethernet interface, use `diagnostic interface ethernet` *interface-number*.
   f. For the rack locator, use `diagnostic rack-locator`
   g. In all cases, the number of the device can be substituted with the keyword `all` to run diagnostics against all such devices in the chassis.
4. While in diagnostic mode, set the test type to LED by executing `test led`.
5. You will be using the default iterations parameter, which is 0. This means that the command will continue running until you type `stop`. If you were already in diagnostics mode before you began this procedure, you should set iterations to 0 by typing `iterations 0`.
6. Start the test by entering `Start`.
7. Go to the FRU and observe the identify LED flashing.
8. Stop the test by entering `stop`.
9. **This procedure ends here.**

Filing bug reports

To initiate bug reports, contact your service provider.

Diagnostics for networks managed by IBM Network Manager

IBM Network Manager diagnostic procedures used to perform diagnose link problems.

Note: IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.

Before you run diagnostics for the GX HCA-based InfiniBand (IB) networks, read through the procedures and then perform the steps outlined in "General link diagnostics procedure."

Link diagnostic procedures

Run diagnostic procedures to help isolate to the correct FRU on a failing link.

Do this procedure using the IBM Network Manager.

General link diagnostics procedure

To isolate to a failing FRU, use this procedure as the general instructions for running link diagnostic procedures.

For a switch-to-switch link, run switch port diagnostics on both ends of the link. For a switch-to-adapter link, run switch port diagnostics on the switch port and GX adapter port diagnostics on the GX adapter.

If you have not determined which devices are at each end of the cable on the link, use one of the procedures in "Procedures for finding FRUs" on page 218.

If you already know which devices are at each end of the link cable, use the following procedure to isolate to a single FRU:

Note: If a link uses an octopus cable to connect a 12x device to three, 4x switch ports, carefully record which switch ports the 4x connectors of the octopus cable connect to. Make sure you replace them in the correct order. If you do not replace them in the correct order, you will need to determine the configuration again using the procedure in "Planning for InfiniBand networks" on page 5. For more on the behavior of 12x connections see "Planning octopus cables in static 12x cabling" on page 19..

Once you have replaced or reseated the cable, verify connectivity using the procedure found in "Verifying static-12x mode connectivity" on page 67.

1. Run Switch Port Link diagnostics on the port that is reporting or indicating the event that requires link diagnostics to be run for FRU isolation. Perform the steps in "Switch port link diagnostic procedure" on page 166.

   Was a faulty FRU indicated?

   No: go to the next step.

   Yes: Replace the faulty FRU and reconnect the cable back on to the link. **This ends this procedure**
2. Run link diagnostics on the port at the other end of the cable. Use the following table to determine which diagnostics to run.

© Copyright IBM Corp. 2004, 2007

Device end of the cable	Procedure
Switch	Perform the procedure in "Switch port link diagnostic procedure."
GX adapter	Perform the procedure in "GX adapter port diagnostic procedure" on page 170.
Other adapter	Refer to the service procedures for the adapter or the server containing the adapter and run any available diagnostics.

Was a faulty FRU indicated?

No: go to the next step.

Yes: Replace the faulty FRU and reconnect the cable back on to the link. **This ends this procedure**

3. Go to the switch end that is involved in the incident under which you were instructed to run link diagnostics. The FRU location code for the cable end is the first cable end listed in the FRU list in Service Focal Point. The part number of the switch end of the cable is IBNCCAB, and the FRU class is Symbolic procedure.
4. Replace the cable on the link.

   Note: If a link uses an octopus cable to connect a 12x device to three, 4x switch ports, carefully record which switch ports the 4x connectors of the octopus cable connect to. Make sure you replace them in the correct order. If you do not replace them in the correct order, you will need to determine the configuration again using the procedure in "Planning for InfiniBand networks" on page 5. For more on the behavior of 12x connections see "Planning octopus cables in static 12x cabling" on page 19..

   Once you have replaced or reseated the cable, verify connectivity using the procedure found in "Verifying static-12x mode connectivity" on page 67.
5. If the cable replacement does not resolve the problem, start replacing FRUs (in the order in which they are presented in the service procedures that led you here). Or, replace the FRUs in the order in which they are presented in the FRU list.
6. **This procedure ends here.**

Switch port link diagnostic procedure

Note: If you are here because an application problem suggests that a link is faulty (rather than because a serviceable event's isolation procedure led you to this procedure), try to ensure that the application that has the problem is running while you perform the following procedures.

Before performing the switch port link diagnostic procedure, do the following:

1. In the IBM Network Manager View Switch Topology window, open the Port Statistics for diagnostics.
2. Observe the port error counters to determine if errors are being recorded.

   Note: Throughout the Switch port link diagnostic procedure you will see references to the word *counter*, the following table identifies the counters and gives a description of each:
3. Try a new cable, or swap cables between two ports, and observe port error counters to determine if errors are being recorded.

   Notes:

   - If a link uses an octopus cable to connect a 12x device to three, 4x switch ports, there is only one port that is active. Use the active port to observe the port statistics counters in this test. To identify the active port according to the switch model and the group of 3 switch ports which comprise the link, see the tables in the procedure "Verifying static-12x mode connectivity" on page 67. For more on the behavior of 12x connections see "Planning octopus cables in static 12x cabling" on page 19.

- If a link uses an octopus cable to connect a 12x device to three, 4x switch ports, carefully record which switch ports the 4x connectors of the octopus cable connect to. Make sure you replace them in the correct order. If you do not replace them in the correct order, you will need to determine the configuration again using the procedure in "Planning for InfiniBand networks" on page 5. For more on the behavior of 12x connections see "Planning octopus cables in static 12x cabling" on page 19.

4. Put the cable back on the port after any recommended FRU replacement.

   Notes:

   - If the switch port has a 4x connector from an octopus cable attached to it, verify connectivity using the procedure found in "Verifying static-12x mode connectivity" on page 67.
   - If a link uses an octopus cable to connect a 12x device to three, 4x switch ports, there is only one port that is active. Use the active port to observe the port statistics counters in this test. To identify the active port according to the switch model and the group of 3 switch ports which comprise the link, see the tables in the procedure "Verifying static-12x mode connectivity" on page 67. For more on the behavior of 12x connections see "Planning octopus cables in static 12x cabling" on page 19.

5. Document the results in Service Focal Point.

Table 22. Port counter field table

Port counter field	Description
Symbol errors	Total number of symbol errors detected on one or more lanes.
Link recovery errors	Total number of times the port training state machine has successfully completed the link error recovery process.
Link downs	Total number of times the port training state machine has failed the link error recovery process and downed the link.
Port received errors	Total number of packets containing an error that were received on the port. These type of errors include: • Local physical errors (ICRC, VCRC, FCCRC, and physical errors that cause entry into bad) • Malformed data packet errors (Lver, length, VL) • Malformed link packet errors (operand, lengh, VL) • Packets discarded due to buffer overrun
Port received remote physical errors	Total number of packets marked with the EBP delimiter received on the port.
Port received switch relay errors	Total number of packets received on the port that were discarded because they could be forwarded by the switch relay. Reasons for this include: • DLID mapping • VL mapping • Looping (output port = input port).
Port transmit discards	Total number of outbound packets discarded by the port because the port is down or congested. Reasons for this include: • Output port is in the inactive state • Packet length exceeded neighbor MTU • Switch lifetime limit exceeds • Switch HOQ limit exceeds

Table 22. Port counter field table (continued)

Port counter field	Description
Port transmit constraint errors	Total number of packets not transmitted from the port for the following reasons: • FilterRawOutbound is true and packet is raw • PartitionEnforcementOutbound is true and packet fails partition key check, IP version check, or transport header version check.
Port received constraint errors	Total number of packets received on the port that are discarded for the following reasons: • FilterRawInbound is true and packet is raw • PartitionEnforcementInbound is true and packet fails partition key check, IP version check, or transport header version check.
Local link integrity errors	The number of times that the frequency of packets containing local physical errors exceeded local_phy_errors.
Excessive buffer overrun errors	The number of times that overrun errors consecutive flow control update periods occurred with at least one overrun error in each period.
VL15 droppeds	Number of incoming VL15 packets dropped due to resource limitations on port selected by PortSelect.
Port transmit data	**Optional:** shall be zero if not implemented. Total number of data octets, divided by 4, transmitted on all VLs from the port selected by PortSelect. This includes all octets between (and not including) the start of packet delimiter and VCRC. It excludes all link packets. • An implementer might choose to count data octets in groups larger than four but are encouraged to choose the smallest group possible. Results are still reported as a multiple of four octets. • The port transmit data packets are counting user data being passed, even though port transmit packets may not be increasing. The random pattern that is being passed on idle cycles, provide enough switching to help diagnose the link.
Port received data	**Optional:** shall be zero if not implemented. Total number of data octets, divided by 4, received on all VLs from the port selected by PortSelect. This includes all octets between (and not including) the start of packet delimiter and VCRC. It excludes all link packets. • An implementer might choose to count data octets in groups larger than four but are encouraged to choose the smallest group possible. Results are still reported as a multiple of four octets. • The port received data packets are counting user data being passed, even though port received packets may not be increasing. The random pattern of data that is being passed on idle cycles, provide enough switching to help diagnose the link.
Port transmit packets	**Optional:** shall be zero if not implemented. Total number of data packets, excluding link packets, transmitted on all VLs from the port selected by PortSelect.
Port received packets	**Optional:** shall be zero if not implemented. Total number of data packets, excluding link packets, received on all VLs from the port selected by PortSelect.

To perform the **Switch port link diagnostic procedure**, do the following:

1. Obtain a spare cable for this test.
2. Log in to the HMC that is running the IBM Network Manager.
3. Click **Switch Management > IBM Network Manager > View Switch Topology**.
4. Select the switch port for which you want to run diagnostics.

Note: If a link uses an octopus cable to connect a 12x device to three, 4x switch ports, there is only one port that is active. Use the active port to observe the port statistics counters in this test. To identify the active port according to the switch model and the group of 3 switch ports which comprise the link, see the tables in the procedure "Verifying static-12x mode connectivity" on page 67. For more on the behavior of 12x connections see "Planning octopus cables in static 12x cabling" on page 19.

5. From the menu, click **Selected > Diagnose > Port Statistics**.
6. Verify that the Port Location Code field indicates the port that you want to test. If it does not indicate the correct port, click **Cancel** and go back to 3 on page 168.
7. Click **Clear Counters**.
8. Wait for several minutes.
9. Click **Refresh**.
10. If the counters are still 0, the port is not logging any errors. Perform the following procedure. If the counters are not 0, proceed to the next step.
    a. Check the LEDs on the switch port(s) on either end of the cable.
        1) If the LEDs are on or flashing, continue to step 10b.
        2) If the LEDs are off, make a note that the LED was off and go to step 11.
    b. If you were directed here from an isolation procedure in a serviceable event, check the duplicate count for the serviceable event in the Serviceable Event Details window:
        1) If the duplicate count is 0, make a comment in the serviceable event that the link diagnostics do not indicate that service should be performed, and close the link.
        2) If the duplicate count is not 0, proceed to step 11.
    c. If you are here because of an application problem and the application was running while you were observing the error counters, this link is not contributing to the problem.
11. Click **Diagnose** and follow the instructions:
    a. Verify that the port location code matches the port your are attempting to diagnose. If it does not match, click **Cancel** and select the correct port.
    b. You are instructed to replace a cable connected to the port with loopback adapter to enable diagnosis. Because there is no loopback adapter currently available, you can either replace the cable with a new cable, or swap the cable for this switch port with a cable that is known to be good on another switch port. If the cable to the switch is swapped, ensure that the connection to the adapter port is also swapped.

        Notes:

        - If a switch port has a 4x connector from an octopus cable attached to it, carefully record which port each of the 4x connectors of the octopus cable connects to. Connecting them incorrectly will cause this test to fail.
        - If you swap the cables connected to this port, you will induce errors on the other port, and might cause application failures. A serviceable event for a link down is reported to Service Focal Point within 5 to 10 minutes of pulling the cables.
        - When unplugging the cable on the port, check for bent pins or an obviously poor connection that might cause the errors that have been occurring on this port.

    c. After you have installed the cable, click **Next**.
        - If you installed the cable on the incorrect port, you will be notified.
        - Note the time.
        - If you put a cable on the incorrect port, make note of the location, because a serviceable event will be created for that port, and you will need to close it out. Instructions for doing so are at the end of the Switch port diagnostics procedure.
    d. Click **Clear Counters**.
    e. Check the counters.

1) If the counters are not 0, click **Refresh**, and wait five minutes.
2) If the counters are 0, check the LEDs on the switch port(s) on either end of the cable.
   a) If the LED is on or flashing and it was off in step10 on page 169, replace the original cable with a good cable.
   b) If the LEDs are on or flashing, go to step 12.
   c) If the LEDs are off and they were also off in step 10 on page 169, replace the switch card.
   d) If the LEDs are off and they were on or flashing in step 10 on page 169, check the cable connection, and retry step 11 on page 169 Try substituting a different cable to see if the LEDs can be returned to an on or flashing state.

f. After clicking the **Refresh** button, wait 5 minutes.
g. If any of the counters advance beyond 0 again, replace the switch card.

   Note: In Service Focal Point, make sure that you document that you replaced the switch card.

h. Click **Next**.
i. You will be asked if you wish to create a serviceable event for tracking purposes.
   1) If you started running diagnostics because of a serviceable event or no problem was found, choose **None** for the device to report to Service Focal Point. In this way, no new serviceable event will be generated.
   2) If you started running diagnostics because an application issue led you to believe that this port has a problem, and there was not an associated serviceable event, you should choose the device to which you isolated the problem.
      a) Go to Service Focal Point and look for the serviceable event. The reference code will be CBFF0000 if you chose the port with the loopback adapter on it. It will be CBFF0001 if you chose the port on the other side of the cable. It will be CBFF0002 if you chose the cable. Note that you will generally only choose the cable or the port on the other side of the cable after you have run loopback diagnostics on the ports on either side of the cable.
      b) Make appropriate comments and indications that you've replaced the FRU in the FRU list.
      c) Close out the event.
j. Click **Next**.
k. Reinstall the cable on to the port.

   Notes:

   - If the switch port has a 4x connector from an octopus cable attached to it, be sure to reattach the 4x connectors to the proper 4x switch ports
   - If the switch port has a 4x connector from an octopus cable attached to it, verify connectivity using the procedure found in "Verifying static-12x mode connectivity" on page 67.

12. If you placed a cable on the incorrect port, a false serviceable event is created. You must close out serviceable events in Service Focal Point, keeping in mind that the reporting device will be the switch, and not the adapter.
    a. Make a comment in the false serviceable event to indicate that it was caused by placing a looped cable on the incorrect port.
    b. Close the false serviceable event.

GX adapter port diagnostic procedure

Notes:

1. Throughout the GX adapter port diagnostic procedure you will see references to the word *counter*. Refer to the Port counter field table in the Switch port link diagnostic procedure for an explanation of the counter types.

2. Before replacing an HCA, review "Installing or replacing a GX Host Channel Adapter" on page 64. If you are considering deferred maintenance of the adapter, review "Deferring replacement of a failing Host Channel Adapter" on page 66.

To perform the GX adapter port diagnostic procedure, do the following:

1. Obtain a spare cable for this test.
2. Log in to the HMC that is running the IBM Network Manager.
3. Go to **Switch Management > IBM Network Manager > View Endpoint Topology.**
4. Select the adapter port for which you want to run diagnostics.
5. From the menu choose **Selected > Diagnose > Port Statistics**.
6. Verify that the Port Location Code field indicates the port that you want to test. If it does not indicate the correct port, click **Cancel** and go back to step 3.
7. Click **Clear Counters**.
8. Wait for several minutes.
9. Click **Refresh**.
10. If the counters are still 0, this indicates that the port is not logging any errors. Perform the following procedure. If the counters are not 0, proceed to the next step.
    a. Check the LEDs on the switch port on the other end of the cable.
        1) If the LED is on or flashing, continue to step 10b.
        2) If the LED is off, make a note that the LED was off and go to step 11.
    b. If you were directed here from an isolation procedure in a serviceable event, check the duplicate count for the serviceable event in the Serviceable Event Details window.
        1) If the duplicate count is 0, you should make a comment in the serviceable event that link diagnostics do not indicate that service should be performed, and close the link.
        2) If the duplicate count is not 0, proceed to step 11.
    c. If you are here because of an application issue and the application was running while you were observing the error counters, this link is not contributing to the issue.
11. Click **Diagnose** and follow the instructions:
    a. Verify that the port location code matches the port your are attempting to diagnose. If it does not match, click **Cancel** and select the correct port.
    b. You are instructed to replace the cable connected to the port with loopback adapter to enable diagnosis. Because there is no loopback adapter currently available, you can either replace the cable with a new cable, or swap the cable for this adapter port with a cable that is known to be good on another adapter port. If the cable to the adapter is swapped, ensure that the connection to the switch port is also swapped.

        Note:

        - If a switch port has a 4x connector from an octopus cable attached to it, carefully record which port each of the 4x connectors of the octopus cable connects to. Connecting them incorrectly will cause this test to fail.
        - If you swap the cables connected to this port, you will induce errors on the other port, and might cause application failures. A serviceable event for a link down is reported to Service Focal Point within 5 to 10 minutes of pulling the cables.
        - When unplugging the cable on the port, check for bent pins or an obviously poor connection that might cause the errors that have been occurring on this port.

    c. After you have installed the cable, click **Next**.
        - If you installed the cable on the incorrect port, you will be notified.
        - Note the time.

- If you put a cable on the incorrect port, make note of the location, because a serviceable event will be created for that port, and you will need to close it out. Instructions for doing so are at the end of the Switch port diagnostics procedure.

d. Click **Clear Counters**.

e. Check the counters:

1) If the counters are not 0, click **Refresh**. Wait five minutes.

2) If the counters are 0, check the LED on the switch port on the other end of the cable from the adapter.

a) If the LED is on or flashing and it was off in step10 on page 171 replace the original cable with a good cable.

b) If the LED is on or flashing, go to step 12 on page 173.

c) If the LED is off and it was also off in step 10 on page 171, replace the adapter.

Note: Before replacing an HCA, review "Installing or replacing a GX Host Channel Adapter" on page 64. If you are considering deferred maintenance of the adapter, review "Deferring replacement of a failing Host Channel Adapter" on page 66.

d) If the LED is off and it was on or flashing in step 10 on page 171, check the cable connection and retry step 11 on page 171. Try substituting a different cable to see if the LEDs can be returned to an on or flashing state.

f. After clicking the **Refresh** button, wait 5 minutes.

g. If any of the counters advance beyond 0, replace the adapter.

Notes:

1) Before replacing an HCA, review "Installing or replacing a GX Host Channel Adapter" on page 64. If you are considering deferred maintenance of the adapter, review "Deferring replacement of a failing Host Channel Adapter" on page 66.

2) In Service Focal Point, make sure that you document that you replaced the adapter.

h. Click **Next**.

i. You will be asked if you want to create a Serviceable Event for tracking purposes.

1) If you started running diagnostics because of a serviceable event or no problem was found, choose **None** for the device to report to Service Focal Point. In this way, no new serviceable event will be generated.

2) If you started running diagnostics because you believed this port to have a problem, and there was not an associated serviceable event, you should choose the device to which you isolated the problem.

a) Go to Service Focal Point and look for the serviceable event. If you chose the port with the cable on it, the reference code will be CBFF0000. If you chose the port on the other side of the cable, the reference code will be CBFF0001. It will be CBFF0002 if you chose the cable. Note that you will generally only choose the cable or the port on the other side of the cable after you have run loopback diagnostics on the ports on either side of the cable.

b) Make appropriate comments and indications that you have replaced the FRU in the FRU list.

c) Close out the event.

j. Click **Next**.

k. Replace the cable back on to the port.

l.

- If the switch port has a 4x connector from an octopus cable attached to it, be sure to reattach the 4x connectors to the proper 4x switch ports
- If the switch port has a 4x connector from an octopus cable attached to it, verify connectivity using the procedure found in "Verifying static-12x mode connectivity" on page 67.

12. If you placed a cable on the incorrect port, a false serviceable event is created. You must clean up Service Focal Point keeping in mind that the reporting device will be the switch, and not the adapter.
    a. Make a comment in the false serviceable event to indicate that it was caused by placing a looped cable on the incorrect port.
    b. Close the false serviceable event.

InfiniBand switch card diagnostics

This procedure describes how to run diagnostics on cards using the InfiniBand switch command-line interface.

In some isolation procedures, you may be instructed to run switch card diagnostics, which are not currently available from the "InfiniBand switch reference information" on page 3. The following procedure describes how to run diagnostics on switch cards using the InfiniBand switch command-line interface.

Note:

- If you need to run diagnostics on a switch card, the card will be temporarily unavailable to perform its normal functions. This may result in downstream errors that are reported to Service Focal Point. Therefore, you should be prepared to close out Service Focal Point events reported by devices connected to the card, and by devices monitoring the card. After running diagnostics, wait 5 to 10 minutes for the diagnostics to flow through the service subsystem and check for serviceable events that might be caused by this action (especially Link Down events), and close them out as appropriate. Also, note in the comments for the false events that they were caused by running card diagnostics
- If you want to turn on or off identification LEDs, you should refer to "Manipulating FRU identification LEDs" on page 136.

1. To access the administrative port for the switch chassis that contains the card, use the procedure in "Accessing a serial port on a switch" on page 157. You will require "admin" privileges.
2. You will need to put the switch into the Global Configuration mode that is described in the Topspin *Command Line Interface Reference Guide* document. To locate the Topspin documentation, go to "InfiniBand switch reference information" on page 3.
    a. In User Exec mode, enter the **enable** command to enter Privileged Exec mode.
    b. Enter the **configure terminal** command to enter Global Configuration mode.
3. In the switch command-line interface, you will use the **diagnostic** command to enter diagnostic mode. For details on the diagnostic command as described in the Topspin *Command Line Interface Reference Guide*, go to the "InfiniBand switch reference information" on page 3. Use one of the following options to choose the appropriate card:
    a. For LIMs, or Fabric Control Module use: **diagnostic card** *slot_number*.
    b. For power supplies (or combined power/fan modules), use: **diagnostics power-supply** *supply*.
    c. For fans, use: **diagnostic fan** *fan-number*.
    d. For the chassis module, use: **diagnostic chassis**.
    e. For the Ethernet interface, use **diagnostic interface ethernet** *interface-number*.
    f. For the rack locator, use **diagnostic rack-locator**.
    g. In all cases, the number of the device can be substituted with the keyword **all** to run diagnostics against all such devices in the chassis.
4. While in diagnostic mode, set the test type to self-test by running **test self-test**.
5. Set the number of iterations to one by executing **iterations 1**.
6. Start the test by executing **start**.
7. Check the results of the test by executing **more** *slot-id***:syslog:hwif_log**

If the switch is a 7048-120 or SFS7000P the *slot-id* is always 1. If the switch is a 7048-270 or SFS7008P, the *slot-id* is either 11 or 12 depending on which switch card is set as the master.

8. If there is a self-test failure, replace the card.
   a. Turn on the identify LED so that you know which card to replace. See "Accessing FRU Identification LEDs" on page 162 for information concerning turning on and off LEDs.
   b. Turn off the iterations parameter so that the LED will flash until you enter `Stop`. On the Command Line Interface (CLI), enter `iterations 0`.
   c. Start the identify LED flashing. On the CLI, enter `start`.
   d. Replace the card.
   e. If the identify LED on the new card is flashing, on the CLI, enter `stop`.
9. **This procedure ends here**.

Review syslog:hwif_log

Results of entering the command `more` *`slot-id`*`:syslog:hwif_log` will look like the following. All of the tests in this example have PASSED.

```
Fri Aug 12 10:05:45 2005: POST: BASEBOARD in Slot 1: (RTC): PASSED
Fri Aug 12 10:05:45 2005: POST: BASEBOARD in Slot 1: (FPGA: Revision = 0xaa): PASSED
Fri Aug 12 10:05:46 2005: POST: BASEBOARD in Slot 1: (SEEPROM DATALOG): PASSED
Fri Aug 12 10:05:46 2005: POST: BASEBOARD in Slot 1: (SUMMIT): PASSED
Fri Aug 12 10:05:46 2005: POST: BASEBOARD in Slot 1: (VPD): PASSED
Fri Aug 12 10:05:46 2005: Set POST Status: slot 1 status 1 error code 1 (0) 2 (0) 3 (0) 4 (0)
Fri Aug 12 10:05:50 2005: POST: BASEBOARD in Slot 1: (FAN 1 Slot 4): PASSED
Fri Aug 12 10:05:50 2005: POST: BASEBOARD in Slot 1: (FAN 2 Slot 4): PASSED
Fri Aug 12 10:05:50 2005: Set POST Status: slot 4 status 1 error code 1 (0) 2 (0) 3 (0) 4 (0)
Fri Aug 12 10:05:58 2005: POST: BASEBOARD in Slot 1: (FAN 3 Slot 5): PASSED
Fri Aug 12 10:05:58 2005: POST: BASEBOARD in Slot 1: (FAN 4 Slot 5): PASSED
Fri Aug 12 10:05:58 2005: Set POST Status: slot 5 status 1 error code 1 (0) 2 (0) 3 (0) 4 (0)
Fri Aug 12 10:05:58 2005: card_startup.x : card is starting up
Fri Aug 12 10:05:58 2005: Reset reason is power-on
Fri Aug 12 10:06:05 2005: POST: Anafa2 firmware check: PASSED
Fri Aug 12 10:06:19 2005: POST: Anafa2 DMA test: PASSED
Fri Aug 12 10:06:31 2005: Reporting POST passed on FRU 1
Fri Aug 12 10:06:31 2005: Reporting POST passed on FRU 4
Fri Aug 12 10:06:31 2005: Reporting POST passed on FRU 5
```

Symbolic FRUs for InfiniBand cluster networks that are managed by the IBM Network Manager

Describes how symbolic FRUs are used to diagnose to a failing FRU.

Note: IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.

To find a FRU, you must first understand the information that you have at hand to identify it. The first key is the FRU location code, which, displays in the FRU list for a serviceable event. Valid switch and adapter FRU location codes are provided below.

If you do not have a valid switch or adapter location code, you will have to perform a few more steps to get to the point where you understand the location of the device.

The following table will help you find which procedure you need to use:

	Desired FRU	
Symptom	**Switch FRU**	**Adapter FRU**
Valid Location Code for Desired FRU	"How to find a switch FRU with a valid location code" on page 222	"How to find an adapter FRU with a valid location code" on page 222
Valid Location for Device to which FRU is connected	"How to find a switch FRU using another device with a valid location code" on page 223	"How to find an adapter FRU using another device with a valid location code" on page 223
Reference Code Extension	"How to find a switch FRU using a reference code extension" on page 224	"How to find an adapter FRU using a reference code extension" on page 176
Using the Error Description	"How to determine a switch FRU using the error description" on page 227	"How to determine an adapter FRU using the error description" on page 227

Symbolic FRU format

All symbolic FRUs begin with an explanation of the conditions under which a symbolic FRU is used, then instructions are provided for finding the missing information for the FRU.

Symbolic FRU procedures are used by error log analysis for the IBM Network Manager for either of the following reasons:

- The part number is unknown, and therefore you need to know how to determine it.
- The location code is unknown, and therefore you need to know how to determine the part's location.

If the location code is valid, then the part number is unknown and must be determined manually. See "How to determine an unknown part number for an InfiniBand switch network component" on page 229.

If the location code is not valid, then the location of the component must be found manually.

© Copyright IBM Corp. 2004, 2007

The following is a list of symbolic FRUs listed in alphabetic order:

"CBLCONT" on page 178
"IBNACRD" on page 178
"IBNAPRT" on page 178
"IBNCCAB" on page 179
"IBNPBAT" on page 181
"IBNPFAN" on page 182
"IBNPPOW" on page 182
"IBNSCHS" on page 183
"IBNSCRD" on page 184
"IBNSCTL" on page 185
"IBNSPLN" on page 186
"IBNSSVC" on page 186

How to find an adapter FRU using a reference code extension

How to find an adapter FRU using a reference code extension.

In some cases, there are no valid location codes in a FRU list. In such cases, you will need to work with a logical location id which is recorded in the reference code extension in SFP. Because the IBM Network Manager only presents serviceable events that come from switches, the Reference Code Extension will always be for the switch attached to the desired adapter FRU. So, the basic approach for this procedure is to first find the switch port that reported the serviceable event, and then find the adapter on the other side of the cable.

Note: For this procedure to work, the network should not have been re-cabled since the time of the latest occurrence of the serviceable event. It is particular important that the failing link was not re-cabled.

1. Record Service Focal Point Information:
   a. Record the location code for the attached switch FRU which is also in the FRU list. Everything except for the U[*MachineType*].[*model*].[*serial number*] portion of the location code should be valid.
   b. Record the reference code extension from the serviceable event.
   c. Record the frame, card number and port number from the reference code extension:

Network	Frame	Chassis	Chassis Type	Device Type	Card	Chip	Port
4 Hex character	3 Hex character	2 Hex character	Hex character	Hex character	2 Hex characters	Hex character	2 Hex characters

   If there are fewer than 16 characters, then leading zeroes have been dropped, for example: 0001002052301100

   network = 1; frame=2; cage=5; cage type=5; device type=3; card=01; chip=1; port=00
2. Find and record information about the adapter:
   a. Go to the IBM Network Manager View Switch Topology window.
   b. Using the frame and location code for the attached switch recorded above, find the attached switch.
   c. Record the Neighbor Location-Code, Neighbor Name and Neighbor Frame:Cage. These are the location code, name and frame and cage of the desired FRU. It is possible that the location code found in the View Switch Topology window is now valid.
3. Turn on the Identify LED
   a. Go to the controlling HMC for the managed system. The managed system should be identified by one of the following methods based on information gathered about the adapter:
      1) If you found a valid location code, the managed server is identified in the unit location field U[*FeatureCode*].001.[*SerialNumber*]-.
      2) If this is a high-end server in a 24-inch frame, the frame and cage numbers will help identify the location of the server.
      3) The name of the chassis can identify the server.
   b. Using Service Focal Point-Service Utilities turn on the Identify LED for the managed system.
4. Go to physical location and verify FRU.

- If you know the physical location of the switch by the Chassis name, the MTMS, or the frame, go to that switch, and verify that the Identify LED is on.
- If you don't know the physical location based on the Chassis name, MTMS, or frame, you will have to walk the floor and look for Identify LEDs that are on and then verify that the switch at which you are looking matches the MTMS for the FRU.

How to find an adapter FRU using a reference code extension

How to find an adapter FRU using a reference code extension.

In some cases, there are no valid location codes in a FRU list. In such cases, you will need to work with a logical location id which is recorded in the reference code extension in SFP. Because the IBM Network Manager only presents serviceable events that come from switches, the Reference Code Extension will always be for the switch attached to the desired adapter FRU. So, the basic approach for this procedure is to first find the switch port that reported the serviceable event, and then find the adapter on the other side of the cable.

Note: For this procedure to work, the network should not have been re-cabled since the time of the latest occurrence of the serviceable event. It is particular important that the failing link was not re-cabled.

1. Record Service Focal Point Information:
   a. Record the location code for the attached switch FRU which is also in the FRU list. Everything except for the U[*MachineType*].[*model*].[*serial number*] portion of the location code should be valid.
   b. Record the reference code extension from the serviceable event.
   c. Record the frame, card number and port number from the reference code extension:

Network	Frame	Chassis	Chassis Type	Device Type	Card	Chip	Port
4 Hex character	3 Hex character	2 Hex character	Hex character	Hex character	2 Hex characters	Hex character	2 Hex characters

   If there are fewer than 16 characters, then leading zeroes have been dropped, for example: 0001002052301100

   network = 1; frame=2; cage=5; cage type=5; device type=3; card=01; chip=1; port=00
2. Find and record information about the adapter:
   a. Go to the IBM Network Manager View Switch Topology window.
   b. Using the frame and location code for the attached switch recorded above, find the attached switch.
   c. Record the Neighbor Location-Code, Neighbor Name and Neighbor Frame:Cage. These are the location code, name and frame and cage of the desired FRU. It is possible that the location code found in the View Switch Topology window is now valid.
3. Turn on the Identify LED
   a. Go to the controlling HMC for the managed system. The managed system should be identified by one of the following methods based on information gathered about the adapter:
      1) If you found a valid location code, the managed server is identified in the unit location field U[*FeatureCode*].001.[*SerialNumber*]-.
      2) If this is a high-end server in a 24-inch frame, the frame and cage numbers will help identify the location of the server.
      3) The name of the chassis can identify the server.
   b. Using Service Focal Point-Service Utilities turn on the Identify LED for the managed system.
4. Go to physical location and verify FRU.

- If you know the physical location of the switch by the Chassis name, the MTMS, or the frame, go to that switch, and verify that the Identify LED is on.
- If you don't know the physical location based on the Chassis name, MTMS, or frame, you will have to walk the floor and look for Identify LEDs that are on and then verify that the switch at which you are looking matches the MTMS for the FRU.

CBLCONT

A symbolic FRU that directs you to procedures for diagnosing a problem with a link cable.

Note: IBM Service is responsible for this FRU.

Refer to the IBM Systems Hardware Information Center section on Symbolic FRU procedures.

IBNACRD

This Symbolic FRU procedure represents an InfiniBand adapter card.

Note: IBM Service is responsible for this FRU.

Symbolic FRU procedures are used by Error Log Analysis for the IBM Network Manager for either of the following reasons:

- The part number is unknown, and therefore you need to know how to determine it.
- The location code is unknown, and therefore you need to know how to determine the part's location.

IBNACRD: Unknown Part Number

Use this procedure if the part number is unknown, but the location is known.

If the location code is valid, then the part number is unknown and must be determined manually. Perform the procedure: "How to determine an unknown part number for an InfiniBand switch network component" on page 229.

IBNACRD: Invalid Location Code

If the location code is not valid, then you must locate the component manually.

Use one of the following procedures:

- If there are no valid location codes in the FRU list, use "How to find an adapter FRU using a reference code extension" on page 176.
- Because the part number has been replaced with this Symbolic FRU procedure, you must determine it manually. Perform the procedure: "How to determine an unknown part number for an InfiniBand switch network component" on page 229.

The procedure ends here.

IBNAPRT

This Symbolic FRU procedure represents an InfiniBand adapter port.

Note: IBM Service is responsible for this FRU.

Symbolic FRU procedures are used by Error Log Analysis for the IBM Network Manager for either of the following reasons:

- The part number is unknown, and therefore you need to know how to determine it.
- The location code is unknown, and therefore you need to know how to determine the part's location.

IBNAPRT: Unknown Part Number

Use this procedure if the part number is unknown, but the location is known. If the location code

is valid, then the part number is unknown and must be determined manually. Perform the procedure: "How to determine an unknown part number for an InfiniBand switch network component" on page 229.

This procedure ends here.

IBNAPRT: Invalid Location Code

If the location code is not valid, then the location of the component must be found manually.

Use one of the following procedures:

1. If the switch side has a valid location code, use "How to find a switch FRU using another device with a valid location code" on page 223.
2. If there are no valid location codes in the FRU list, use "How to find an adapter FRU using another device with a valid location code" on page 223.
3. Because the part number has been replaced with this Symbolic FRU procedure, you must determine it manually. Perform the procedure: "How to determine an unknown part number for an InfiniBand switch network component" on page 229.
4. **The procedure ends here**.

IBNCCAB

This symbolic FRU procedure represents an InfiniBand switch cable end.

Note: IBM Service is responsible for this FRU.

Cable ends are represented with a symbolic FRU. Because there is no electronic method for determining cable length and part number, you must look at the cable to determine the part number.

Symbolic FRU procedures are used by Error log analysis for the IBM Network Manager any of three reasons:

- The part number is unknown, and therefore you need to know how to determine it.
- The location code is unknown, and therefore you need to know how to determine the part's location.
- The location code is known, but the switch port is not in a recognized switch machine type.

Note: If this is an octopus cable which connects a 12x device to three, 4x switch ports, only the location of the active switch port is reported for this cable. When replacing or reseating this cable, you must carefully record to which switch ports the 4x connectors of the octopus cable connect. Make sure you replace them in the correct order. If you do not replace them in the correct order, you will need to determine the configuration again using the procedure in "Planning for InfiniBand networks" on page 5. For more information on 12x connections see "Planning for InfiniBand networks" on page 5.

To identify the active port according to the switch model and the group of 3 switch ports which comprise the link, see the tables in the procedure "Verifying static-12x mode connectivity" on page 67. For more on the behavior of 12x connections see "Planning octopus cables in static 12x cabling" on page 19.

Once you have replaced or reseated the cable, verify connectivity using the procedure found in "Verifying static-12x mode connectivity" on page 67.

IBNCCAB: Unknown Part Number

If the location code is valid, then the part number is unknown and must be determined manually. Look at the cable and reference it to the parts catalog to get the part number.

IBNCCAB: Invalid Location Code

If the location code is not valid, the location of the component must be found manually. Perform the following procedure:

Note: If the other side of the cable is known, it is described in the next FRU in the FRU list, that has the part number that references the symbolic FRU procedure CBLCONT.

1. Record the location code for both the IBNCCAB FRU and the following CBLCONT FRU.
2. Record the Reference Code Extension field.
3. If the entire location code for IBNCCAB is not valid, and you cannot determine the location of the port within the chassis, this is the adapter end of a cable. Go to the FRU labeled CBLCONT and get that location code. If the CBLCONT location code is valid, skip to the *Get the other end of the cable* step; otherwise proceed to the next step.
4. Determine which side is described by the Reference Code Extension field: Now, only the unit location (U*MachineType*.*Model*.*SerialNumber*) should be invalid. The rest of the location code is valid information, and looks like U####.###.#######-P*planar number*-C*card number*-T*port number*. Proceed.
   a. From the Reference Code Extension field, record the frame, chassis, card and port fields.

Network	Frame	Cage	Cage type	Device type	Card	Chip	Port 2
4 Hex character	3 Hex character	2 Hex character	Hex character	Hex character	2 Hex characters	Hex character	Hex characters

   If there are fewer than 16 characters, then leading zeroes have been dropped.
   Example: 0001002052301100
   network = 1; frame=2; cage=5; cage type=5; device type=3; card=01; chip=1; port=00
   b. Compare the IBNCCAB location code, card, and port numbers with the reference code extension card and port numbers. If the IBNCCAB and CBLCONT location code, backplane, card, and port numbers are identical, then you cannot discern which one is described by the reference code extension. Keep the location code of the IBNCCAB.
   c. If the numbers in the location code and the reference code extension match, then the reference code extension describes the IBNCCAB location.
      1) If you determined that the reference code extension matches the IBNCCAB location information, record the IBNCCAB location code as the desired location code.
      2) Otherwise, the reference code extension describes the CBLCONT location. Record the CBLCONT location code as the desired location code.
5. Do the following to identify the cable end:
   a. In the IBM Network Manager, open the View Switch Topology window.
   b. Find and select the device that has the frame number from the reference code extension and what you recorded as the desired location code. The frame number helps you identify the chassis, and the valid portion of the location code identifies the port in the chassis to which the cable end is connected.
6. Identify the other end of the cable:

   Record the Neighbor Location-Code, Neighbor Name, and Neighbor Frame:Cage for the other end of the cable.
7. Turn on the Identify LEDs:
   a. For the cable ends connected to switches, perform the following:
      1) In the View Switch Topology window, find the desired FRU using the location code, chassis name, and frame recorded in previous steps.
      2) From the menu in the View Switch Topology window, select **Selected-Identify-On**.
   b. For the cable end connected to an adapter, if applicable, perform the following:
      1) Go to the controlling HMC for the system unit. The system unit should be identified by one of the following methods based on information gathered about the adapter:
         a) If you found a valid location code, the system unit is identified in the unit location field

U*feature_code*.001.*serial_number*-

.

b) If this is a high-end system unit in a 24-inch frame, and if they are set, the frame and cage numbers will help identify the location of the system unit.

c) The name of the chassis can identify the system unit.

2) Using Service Focal Point-Service Utilities turn on the Identify LED for the system unit.

8. Go to physical location and verify FRU
   a. Based on the devices frame number, chassis number, chassis name, MTMS or feature code and serial number, find the device on the floor and verify that the Identify LEDs are on.
   b. If you cannot identify the device's physical location from the recorded data, you must check the hardware and look for Identify LEDs that are on. Verify that the device at which you are looking matches the information that you have for the FRU.
9. **This procedure ends here**.

IBNCCAB: Not a recognized InfiniBand switch machine type

If the FRU is not in a recognized InfiniBand switch chassis, the service subsystem might not interpret events properly for it, and there may also be warranty issues with the switch. Perform the "IBNNURM" on page 102 isolation procedure.

IBNPBAT

This Symbolic FRU procedure represents an InfiniBand switch battery.

Note: IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.

Symbolic FRU procedures are used by Error Log Analysis for the IBM Network Manager for one of three reasons:

- The part number is unknown, and therefore you need to know how to determine it.
- The location code is unknown, and therefore you need to know how to determine the part's location.
- The location code is known, but the switch is not a recognized InfiniBand switch.

IBNPBAT: Unknown Part Number

Use this procedure if the part number is unknown, but the location is known.

1. If the location code is valid, then the part number is unknown and must be determined manually. Perform the procedure: "How to determine an unknown part number for an InfiniBand switch network component" on page 229.
2. **This procedure ends here.**

IBNPBAT: Invalid Location Code

If the location code is not valid, then the location of the component must be found manually. Most of the location code should be valid, except for the unit location (U*MachineType*.*Model*.*SerialNumber* or U*ProductName*.*SerialNumber*; where *ProductName* is SFS7000P* or SFS7008P*).

1. Perform the procedure: "How to determine a switch FRU using the error description" on page 227, keeping in mind that the only FRUs listed are in the same chassis with the battery, and there are no port-level or cable FRUs in the FRU list. Also recall that the battery is located on a card.
2. Because the part number has been replaced with this Symbolic FRU procedure, you must determine it manually. Perform the procedure: "How to determine an unknown part number for an InfiniBand switch network component" on page 229.

3. **The procedure ends here**.

IBNPBAT: Not a recognized InfiniBand switch machine type

If the FRU is not in a recognized InfiniBand switch chassis, the service subsystem might not interpret events properly for it, and there may also be warranty issues with the switch. Perform the "IBNNURM" on page 102 isolation procedure.

IBNPFAN

This Symbolic FRU procedure represents an InfiniBand (IB) switch fan card.

Note: IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.

Symbolic FRU procedures are used by Error Log Analysis for the IBM Network Manager for any of three reasons:

- The part number is unknown, and therefore you need to know how to determine it.
- The location code is unknown, and therefore you need to know how to determine the part's location.
- The location code is known, but the switch is not a recognized InfiniBand switch.

IBNPFAN: Unknown Part Number

Use this procedure if the part number is unknown, but the location is known.

1. If the location code is valid, then the part number is unknown and must be determined manually. Perform the procedure: "How to determine an unknown part number for an InfiniBand switch network component" on page 229.
2. **This procedure ends here.**

IBNPFAN: Invalid Location Code

If the location code is not valid, then the location of the component must be found manually. Most of the location code should be valid, except for the unit location (`U`*`MachineType`*`.`*`Model`*`.`*`SerialNumber`* or `U`*`ProductName`*`.`*`SerialNumber`*; where *`ProductName`* is SFS7000P* or SFS7008P*).

1. Perform the procedure: "How to determine a switch FRU using the error description" on page 227, keeping in mind that the only FRUs listed are in the same chassis with the fan, and there are no port-level or cable FRUs in the FRU list.
2. Because the part number has been replaced with this Symbolic FRU procedure, you must determine it manually. Perform the procedure: "How to determine an unknown part number for an InfiniBand switch network component" on page 229.
3. **The procedure ends here**.

IBNPFAN: Not a recognized InfiniBand switch machine type

If the FRU is not in a recognized InfiniBand switch chassis, the service subsystem might not interpret events properly for it, and there may also be warranty issues with the switch. Perform the "IBNNURM" on page 102 isolation procedure.

IBNPPOW

This Symbolic FRU procedure represents an InfiniBand switch power card.

Note: IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.

Symbolic FRU procedures are used by Error Log Analysis for the IBM Network Manager for one of three reasons:

- The part number is unknown, and therefore you need to know how to determine it.
- The location code is unknown, and therefore you need to know how to determine the part's location.
- The location code is known, but the switch is not a recognized InfiniBand switch.

IBNPPOW: Unknown Part Number

Use this procedure if the part number is unknown, but the location is known.

1. If the location code is valid, then the part number is unknown and must be determined manually. Perform the procedure: "How to determine an unknown part number for an InfiniBand switch network component" on page 229.
2. **This procedure ends here.**

IBNPPOW: Invalid Location Code

If the location code is not valid, then the location of the component must be found manually. Most of the location code should be valid, except for the unit location (U*MachineType*.*Model*.*SerialNumber* or U*ProductName*.*SerialNumber*; where *ProductName* is SFS7000P* or SFS7008P*).

1. Perform the procedure: "How to determine a switch FRU using the error description" on page 227, keeping in mind that the only FRUs listed are in the same chassis with the power supply, and there are no port-level or cable FRUs in the FRU list.
2. Because the part number has been replaced with this Symbolic FRU procedure, you must determine it manually. Perform the procedure: "How to determine an unknown part number for an InfiniBand switch network component" on page 229.
3. **The procedure ends here**.

IBNPPOW: Not a recognized InfiniBand switch machine type

If the FRU is not in a recognized InfiniBand switch chassis, the service subsystem might not interpret events properly for it, and there may also be warranty issues with the switch. Perform the "IBNNURM" on page 102 isolation procedure.

IBNSCHS

This Symbolic FRU procedure represents an InfiniBand Switch chassis.

Note: IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.

Symbolic FRU procedures are used by Error Log Analysis for the IBM Network Manager for one of three reasons:

- The part number is unknown, and therefore you need to know how to determine it.
- The location code is unknown, and therefore you need to know how to determine the part's location.
- The location code is known, but the switch is not a recognized InfiniBand switch.

IBNSCHS: Unknown Part Number

Use this procedure if the part number is unknown, but the location is known.

1. If the location code is valid, then the part number is unknown and must be determined manually. Perform the procedure: "How to determine an unknown part number for an InfiniBand switch network component" on page 229. You should be able to find the component using "How to find a switch FRU using another device with a valid location code" on page 223.
2. **This procedure ends here.**

IBNSCHS: Invalid Location Code

If the location code is not valid, then the location of the component must be found manually. Most of the location code should be valid, except for the unit location (*UMachineType.Model.SerialNumber* or *UProductName.SerialNumber*; where *ProductName* is SFS7000P* or SFS7008P*). Perform the following procedure:

1. If there is a switch device in the FRU list with a valid location code, use that location code to determine the chassis in which the switch backplane resides. If this is not possible, proceed to the next step.
2. Use the procedure: "How to determine a switch FRU using the error description" on page 227.
3. Because the part number has been replaced with this Symbolic FRU procedure, you must determine it manually. Perform the procedure: "How to determine an unknown part number for an InfiniBand switch network component" on page 229.
4. **The procedure ends here**.

IBNSCHS: Not a recognized InfiniBand switch machine type

If the FRU is not in a recognized InfiniBand switch chassis, the service subsystem might not interpret events properly for it, and there may also be warranty issues with the switch. Perform the "IBNNURM" on page 102 isolation procedure.

IBNSCRD

This Symbolic FRU procedure represents an InfiniBand (IB) Switch Card.

Note: IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.

Symbolic FRU procedures are used by Error Log Analysis for the IBM Network Manager for one of the following reasons:

- The part number is unknown, and therefore you need to know how to determine it.
- The location code is unknown, and therefore you need to know how to determine the part's location. Valid card location codes have the format *Umachine type.model.serial number*-P*x*-C*y*; where *x* is the backplane number and *y* is the card slot number.
- The location code is known, but the switch is not a recognized InfiniBand switch.

IBNSCRD: Unknown Part Number

Use this procedure if the part number is unknown, but the location is known.

1. If the Location Code is valid, then the part number is unknown and must be determined manually. Perform the procedure: "How to determine an unknown part number for an InfiniBand switch network component" on page 229. Find the component using "How to find a switch FRU using another device with a valid location code" on page 223.
2. **This procedure ends here.**

IBNSCRD: Invalid Location Code

If the location code is not valid, then the location of the component must be found manually. Most of the location code should be valid, except for the unit location (*UMachineType.Model.SerialNumber* or *UProductName.SerialNumber*; where *ProductName* is SFS7000P* or SFS7008P*). Perform one of the following procedures:

- If another card in the FRU list has a valid location code, perform "How to find a switch FRU using another device with a valid location code" on page 223
- If there are no valid location codes in the FRU list, perform "How to find a switch FRU using a reference code extension" on page 224.

- Because the part number has been replaced with this Symbolic FRU procedure, you must determine it manually. Perform the procedure: "How to determine an unknown part number for an InfiniBand switch network component" on page 229
- **This procedure ends here.**

IBNSCRD: Not a recognized InfiniBand switch machine type

If the FRU is not in a recognized InfiniBand switch chassis, the service subsystem might not interpret events properly for it, and there may also be warranty issues with the switch. Perform the "IBNNURM" on page 102 isolation procedure.

IBNSCTL

This Symbolic FRU procedure represents an InfiniBand (IB) switch control card.

Note: IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.

Note: After repair procedures, verify network function according to the "Repair verification" on page 150 procedure.

Symbolic FRU procedures are used by Error Log Analysis for the IBM Network Manager for one of three reasons:

- The part number is unknown, and therefore you need to know how to determine it.
- The location code is unknown, and therefore you need to know how to determine the part's location.
- The location code is known, but the switch is not a recognized InfiniBand switch.

IBNSCTL: Unknown Part Number

Use this procedure if the part number is unknown, but the location is known.

1. If the location code is valid, then the part number is unknown and must be determined manually. Perform the procedure: "How to determine an unknown part number for an InfiniBand switch network component" on page 229.
2. **This procedure ends here.**

IBNSCHS: Invalid Location Code

If the location code is not valid, then the location of the component must be found manually. Most of the location code should be valid, except for the unit location (U*MachineType*.*Model*.*SerialNumber* or U*ProductName*.*SerialNumber*; where *ProductName* is SFS7000P* or SFS7008P*).

1. Perform the procedure: "How to determine a switch FRU using the error description" on page 227, keeping in mind that the only FRUs listed are in the same chassis with the controller logic. This is true even if there are port-level FRUs in the FRU list, which refer to the service network connection.
2. Because the part number has been replaced with this Symbolic FRU procedure, you must determine it manually. Perform the procedure: "How to determine an unknown part number for an InfiniBand switch network component" on page 229.
3. **The procedure ends here**.

IBNSCTL: Not a recognized InfiniBand switch machine type

If the FRU is not in a recognized InfiniBand switch chassis, the service subsystem might not interpret events properly for it, and there may also be warranty issues with the switch. Perform the "IBNNURM" on page 102 isolation procedure.

IBNSPLN

This Symbolic FRU procedure represents an InfiniBand switch backplane.

Note: IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.

Symbolic FRU procedures are used by Error Log Analysis for the IBM Network Manager for one of three reasons:

- The part number is unknown, and therefore you need to know how to determine it.
- The location code is unknown, and therefore you need to know how to determine the part's location. Valid planar location codes have the format *Umachine type.model.serial number*-P*x*; where x is the backplane number.
- The location code is known, but the switch is not a recognized InfiniBand switch.

IBNSPLN: Unknown Part Number

Use this procedure if the part number is unknown, but the location is known.

1. If the location code is valid, then the part number is unknown and must be determined manually. Perform the procedure: "How to determine an unknown part number for an InfiniBand switch network component" on page 229. You should be able to find the component using "How to find a switch FRU using another device with a valid location code" on page 223.
2. **This procedure ends here.**

IBNSPLN: Invalid Location Code

If the location code is not valid, then the location of the component must be found manually. Most of the location code should be valid, except for the unit location (U*MachineType.Model.SerialNumber* or U*ProductName.SerialNumber*; where *ProductName* is SFS7000P* or SFS7008P*). Perform the following procedure:

1. If there are no port-level or cable FRUs, which contain –T# in their location codes, perform the following substeps. Otherwise skip to the next.
   a. If there is a switch device in the FRU list with a valid location code, use that location code to determine the chassis in which the switch backplane resides. If this is not possible, proceed to the next step.
   b. Perform the procedure "How to determine a switch FRU using the error description" on page 227.
2. If there are no port-level or cable FRUs, which contain –T# in their location codes, perform the procedure "How to find a switch FRU using a reference code extension" on page 224.
3. Because the part number has been replaced with this Symbolic FRU procedure, you must determine it manually. Perform the procedure: "How to determine an unknown part number for an InfiniBand switch network component" on page 229.
4. **The procedure ends here**.

IBNSPLN: Not a 7048 switch machine type

If the FRU is not in a recognized InfiniBand switch chassis, the service subsystem might not interpret events properly for it, and there may also be warranty issues with the switch. Perform the "IBNNURM" on page 102 isolation procedure.

IBNSSVC

This Symbolic FRU procedure represents an InfiniBand switch service Ethernet connection.

Note: IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.

Symbolic FRU procedures are used by Error Log Analysis for the IBM Network Manager for one of three reasons:

- The part number is unknown, and therefore you need to know how to determine it.
- The location code is unknown, and therefore you need to know how to determine the part's location.
- The location code is known, but the switch is not a recognized InfiniBand switch.

IBNSSVC: Unknown Part Number

Use this procedure if the part number is unknown, but the location is known.

1. If the location code is valid, then the part number is unknown and must be determined manually. Perform the procedure: "How to determine an unknown part number for an InfiniBand switch network component" on page 229.
2. **This procedure ends here.**

IBNSSVC: Invalid Location Code

If the location code is not valid, then the location of the component must be found manually. Most of the location code should be valid, except for the unit location (U*MachineType*.*Model*.*SerialNumber* or U*ProductName*.*SerialNumber*; where *ProductName* is SFS7000P* or SFS7008P*).

1. Perform the procedure: "How to determine a switch FRU using the error description" on page 227, keeping in mind that the only FRUs listed are in the same chassis with the service network connection. The service network connection should be the only port-level FRU in the FRU list.
2. Because the part number has been replaced with this Symbolic FRU procedure, you must determine it manually. Perform the procedure: "How to determine an unknown part number for an InfiniBand switch network component" on page 229.
3. **The procedure ends here**.

IBNSSVC: Not a recognized InfiniBand switch machine type

If the FRU is not in a recognized InfiniBand switch chassis, the service subsystem might not interpret events properly for it, and there may also be warranty issues with the switch. Perform the "IBNNURM" on page 102 isolation procedure.

Status procedures for the IBM Network Manager

The status of your network is graphically displayed by the IBM Network Manager.

Note: For switches other than 7048-120 or 7048-270, the IBM service representative's responsibilities for repair and isolation actions stop at the InfiniBand switch. After the IBM service representative has isolated a problem to a switch or a component within a switch chassis, the responsibility of the service representative is to tell the customer to contact their switch vendor's service organization.

This section is organized by the topology windows that are available for you to view your network. The following table can be used to cross-reference the appropriate window to the subsection that describes the procedures for the status that you are seeing.

Window	Reference
Management Properties View Status Procedures	"Management properties view status procedure"
End-point Topology	"End-Point Topology window status" on page 190
Logical Topology	"Logical topology window status" on page 199
Switch Topology	"Switch topology window status" on page 200
Switch Topology-Switch Properties System Tab	"Switch Topology - Switch Properties: System Status" on page 202
Switch Topology-Port Properties System Tab	"Switch Topology – Port Properties: System Status" on page 203
Switch Environmental Status	"Switch Environmental Status" on page 205

Note:

- The location code for a device is located on the same line as the status.
- The location code is based on the device that is described by that line in the window. The highest level is the switch chassis that contains cards. The next level is the card that contains multiple ports, and finally a single port.

Management properties view status procedure

The Management Properties View provides insight into the switches and IBM Network Manager configuration.

For an overview of the Management Properties view, see "Viewing IBM Network Manager properties" on page 74.

Note: For switches other than 7048-120 or 7048-270, the IBM service representative's responsibilities for repair and isolation actions stop at the InfiniBand switch. After the IBM service representative has isolated a problem to a switch or a component within a switch chassis, the responsibility of the service representative is to tell the customer to contact their switch vendor's service organization.

There is only one status field in this view, and it is under the **Switch** tab. It is the Connectivity column. If it is anything other than **Responsive**, then you should use the Isolation Procedure "IBNSSLC" on page 116.

© Copyright IBM Corp. 2004, 2007

Note: On the bottom left of the window you will see the timestamp for the last auto-update. This represents the last time at which IBM Network Manager updated the information in the Management Properties view.

End-Point Topology window status

Adapter and port status procedures associated with the End-Point Topology window.

For an overview of the End-Point Toplogy window, see "Viewing server topology information in an InfiniBand network" on page 72.

The status fields in the End-point Topology window provide adapter status, adapter power status, and port status, which are described in the following sections:

Note: For switches other than 7048-120 or 7048-270, the IBM service representative's responsibilities for repair and isolation actions stop at the InfiniBand switch. After the IBM service representative has isolated a problem to a switch or a component within a switch chassis, the responsibility of the service representative is to tell the customer to contact their switch vendor's service organization.

Note:

1. Servers are displayed with their Machine Type Model and Serial (MTMS) in the location code column. HCAs and their ports will have IBM location codes in the location code column.
2. You might see servers with PCI HCAs. Because the IBM Network Manager does not gather information about PCI HCAs, you cannot expand the PCI HCAs to get their status or the status of their ports.
3. It is possible for the menu option **Selected-Expand/Collapse** and the twistie indicator in the window to get out of synch, where one may act as if the selected object is expanded, while the other may act as if the selected object is collapsed. If this happens, select another object, then reselect the object of interest.

On the bottom left of the window the timestamp for the last auto-update is displayed. This represents the last time at which IBM Network Manager updated the information in the End-point view.

Description	Reference
Adapter status in the End-Point Topology window supplies a cross-reference table that is used to help identify the adapter status and suggest an appropriate service procedure.	"Adapter status in End-Point Topology window"
Port status in the End-Point Topology window supplies a cross-reference table that is used to help identify the port status and cross-reference to the appropriate procedure for that status.	"Port status in End-Point Topology window" on page 191
Adapter card power status supplies a cross-reference table that is used to help identify the state in the Power status in the End-point topology window and cross-reference to the appropriate procedure for that status.	"Adapter card power status in the End-Point Topology window" on page 191

Adapter status in End-Point Topology window

Use the following table to look up the state in the Adapter Status in the End-point Topology window and cross-reference to the appropriate procedure.

State	Description	Procedure
Not Present	The device is no longer present.	"Adapter status or logical topology state is Not Present" on page 212
Guarded	The device has been guarded and is not available for use. This is likely caused by an internal adapter error, or bus error.	"Adapter status or logical topology state is Garded" on page 212
Functional	The device is present and functional.	This does not imply that the link is available for use, see "Port status in End-Point Topology window" for link information.

Port status in End-Point Topology window

Use the following table to lookup the state in the Port Status in the Endpoint Topology window and cross-reference to the appropriate procedure for that status.

State	Description	Procedure
No Port State Change	The state has not been read by the IBM Network Manager.	"Port status is No Port State Change" on page 213
Down	The link is down.	Procedures for link state as seen by a port for which you will find procedures in the following subsections are logical link states, as opposed to the more physical link training states described in another section. The following table describes in which windows you might see these states. Window Field End-point Topology, Port Status Switch Topology, Port Status, Port Status "Port status is Down" on page 213
Initialized	The link is being initialized.	"Port status is Initialized" on page 213
Armed	The link is changing from initialized to either active or down.	"Port status is Armed" on page 213
Active	The link is available for use.	"Port status is Active" on page 213
Active Defer	The link is status is changing from active to down.	"Port status is Active Defer" on page 214
Reserved Port State	This is not a valid state at this time.	"Port status is Reserved Port State" on page 214

Adapter card power status in the End-Point Topology window

The adapter card power status is seen in the End-Point Topology window. It is relative to an adapter card; therefore, it has no relevance at the port level.

Use the following table to lookup the Power for the adapter card status in the End-Point Topology Window and cross-reference to the appropriate procedure for that status.

State	Description	Procedure
Power Off	There is no power applied to the adapter card.	"Adapter card power status is Power Off"
Power Transition	This indicates that the power state of the adapter card is in the process of transition. This could be from on to off or off to on.	"Adapter card power status is Power Transition"
Power Standby	The power state of the adapter card is at standby.	"Adapter card power status is Power Standby" on page 193
Power IPL Transition	The power to the adapter card has made the transition to a state where the adapter in now IPLing.	"Adapter card power status is Power IPL Transition" on page 193
Up	Power is applied to the adapter card.	"Adapter card power status is Up" on page 193
Power Dump	This indicates a power dump is occurring.	"Adapter card power status is Power Dump" on page 194
Power Termination	The power of the adapter card is starting to transition from on to off.	"Adapter card power status is Power Termination" on page 194
Power Unknown	The state of the adapter card's power is unknown.	"Adapter card power status is Power Unknown" on page 194

Adapter card power status is Power Off

The adapter card power status of Power Off indicates that the power is not applied to the adapter card.

1. If the power should not be applied to the adapter card, this is good status.
2. If power should be applied to the adapter card, perform the following:
   a. Record the adapter card location code from the state has not yet been read by the IBM Network Manager window.
   b. Open Service Focal Point on the HMC that is running the state has not yet been read by the IBM Network Manager.
   c. Look for a serviceable event that relates to power on the adapter card or the server in which the adapter card is installed. Perform the prescribed service for any serviceable event related to that adapter card or the server in which the adapter card is installed.
   d. If there is no serviceable event that relates to power on that adapter card or the server in which the adapter card resides, go to the switch and check the LEDs on the power cards for that switch; refer to "Interpreting LEDs" on page 122.
3. **This procedure ends here.**

Adapter card power status is Power Transition

The adapter card power status of Power Transition indicates that the power state is changing for an adapter card.

1. Wait a few minutes and check the power status again. You should see the Last Auto-update in the bottom left of the window, advance in time.
2. If the adapter card power status of Power Transition is persistent, perform the following procedure:
   a. Record the adapter card location code from the state has not yet been read by the IBM Network Manager window.
   b. Open Service Focal Point on the HMC that is running the state has not yet been read by the IBM Network Manager.

c. Look for a serviceable event that relates to power on that adapter card or the server in which the adapter card is installed. Perform the prescribed service for any serviceable event related to that adapter card or the server in which the adapter card is installed.

d. If there is no serviceable event that relates to power on that adapter card or the server in which the adapter card resides, go to the switch and check the LEDs on the power cards for that switch; refer to "Interpreting LEDs" on page 122.

3. **This procedure ends here.**

Adapter card power status is Power Standby

The adapter card power status of Power Standby indicates that the power is at standby to an adapter card.

1. If the power should not be applied to the adapter card, this is good status.
2. If power should be applied to the switch card, perform the following:
   a. Record the adapter card location code from the state has not yet been read by the IBM Network Manager window.
   b. Open Service Focal Point on the HMC that is running the state has not yet been read by the IBM Network Manager.
   c. Look for a serviceable event that relates to power on that adapter card or the server in which the adapter card is installed. Perform the prescribed service for any serviceable event related to that adapter card or the server in which the adapter card is installed.
   d. If there is no serviceable event that relates to power on that adapter card or the server in which the adapter card resides, go to the switch and check the LEDs on the power cards for that switch; refer to "Interpreting LEDs" on page 122.
3. **This procedure ends here.**

Adapter card power status is Power IPL Transition

The adapter card power status of Power IPL Transition indicates that the power to an adapter card is in a state that has caused the adapter to begin to IPL.

1. If the power should not be applied to the adapter card, this is good status.
2. Wait several minutes and check the status again. You should see the Last Auto-update in the bottom left of the window, advance in time.
3. If the Power Status of Power IPL Transition persists, perform the following:
   a. Record the adapter card location code from the state has not yet been read by the IBM Network Manager window.
   b. Open Service Focal Point on the HMC that is running the state has not yet been read by the IBM Network Manager.
   c. Look for a serviceable event that relates to power on that adapter card or the server in which the adapter card is installed. Perform the prescribed service for any serviceable event related to that adapter card or the server in which the adapter card is installed.
   d. If there is no serviceable event that relates to power on that adapter card or the server in which the adapter resides, go to the switch and check the LEDs on the power cards for that switch; refer to "Interpreting LEDs" on page 122.
4. **This procedure ends here.**

Adapter card power status is Up

The adapter card power status of Up indicates that power is applied to an adapter card. Generally this is good status.

1. If the power should not be applied to the adapter card, perform the following.

a. Power off this adapter card.
b. Check the status of the other adapter cards to see if an adapter card has a power status of Power Off when it should have a power status of Up.
c. If there is an adapter card which has a power status of Power Off when it should have a power status of Up, it could be that the adapter card that has a power status of Power Off was accidentally powered-off instead of the adapter card that has a power status of Up. On the adapter card that has a power status of Power Off, perform the procedure in section "Adapter card power status is Power Off" on page 192.

2. **This procedure ends here.**

Adapter card power status is Power Dump

This status indicates a power issue with an adapter card.

1. If the power should not be applied to the adapter card, this is good status.
2. Wait several minutes and check the status again. You should see the Last Auto-update in the bottom left of the window, advance in time.
3. If the power status of Power Dump persists, perform the following:
   a. Record the adapter card location code from the IBM Network Manager window.
   b. Open Service Focal Point on the HMC that is running the IBM Network Manager.
   c. Look for a serviceable event that relates to power on that adapter card or the server in which the adapter card is installed. Perform the prescribed service for any serviceable event related to that adapter card or the server in which the adapter card is installed.
   d. If there is no serviceable event that relates to power on that adapter card or the server in which the adapter card resides, go to the switch and check the LEDs on the power cards for that switch; refer to "Interpreting LEDs" on page 122.
4. **This procedure ends here.**

Adapter card power status is Power Termination

The adapter card power status of Power Termination indicates that the adapter card status is changing from Up to Power Off, or to Power Standby. This should be a temporary state.

1. If the adapter card was not powered off by a service person, or the customer, wait a few minutes and then check the status again. You should see the Last Auto-update in the bottom left of the window, advance in time.
2. If the status is persistently Power Termination, perform the following procedure:
   a. Record the adapter card location code from the state has not yet been read by the IBM Network Manager window.
3. **This procedure ends here.**

Adapter card power status is Power Unknown

The adapter card power status of Power Unknown indicates that the IBM Network Manager does not currently know the power status of the adapter card. This could be seen at install or after a repair during the period when IBM Network Manager is first acquiring information about the status of the adapter card, or this could be the result of a problem with communicating with the switch.

1. Wait a few minutes and check the status again. You should see the Last Auto-update in the bottom left of the window, advance in time.
2. If the adapter card power status is persistently Power Unknown, perform the following procedure:
   a. Record the location code for the adapter card that has the Power Unknown status from the state has not yet been read by the IBM Network Manager window. Also, note the portion of the location code that indicates the server information: `U[MachineType].[Model].[SerialNumber]`

b. Go to Service Focal Point on the HMC that is running the IBM Network Manager.
c. Look for any serviceable event that was created with either the adapter card, or the backplane, or another card in the FRU list. Use the location code for the adapter card and the location code for the server to do this. The other adapter cards in the same server will have a location code that begins with the location for the server and then continue with the adapter card specific information.
d. If there is a serviceable event against an adapter card in the server, perform the prescribed service.
   1) After performing the prescribed service, recheck the power status.
   2) If the status is still Power Unknown, then continue with the next step. Otherwise, this procedure ends here.
e. Check the Ethernet connections and network between the adapter card and the HMC that is running the state has not yet been read by the IBM Network Manager, and make any necessary repairs.
f. If the status is still Power Unknown, call your next level of service.

3. **This procedure ends here.**

Power status in End-Point Topology window

Power status, relative to the servers power state is seen in the End-point Topology window.

The power status, as seen in the End-point Topology window is relative to a server's power state; therefore, it has no direct relevance at the port and card levels. When the server is in a power state other than Up, the HCA cards and ports are not functional.

Use the following table to look up the state for the Power Status in the End-point Topology window and cross-reference to the appropriate procedure for that status.

State	Description	Procedure
Power Off	There is no power applied to the server.	"Server power status is Power Off"
Power Transition	This indicates that the power state of the server is undergoing a transition. This could be from on to off or off to on.	"Server power status is Power Transition" on page 196
Power Standby	The power state of the server is at standby.	"Server power status is Power Standby" on page 196
Power IPL Transition	The power of the server has undergone a transition to a state where the server is now IPLing.	"Server power status is Power IPL Transition" on page 197
Up	Power is applied to the server.	"Server power status is Up" on page 197
Power Dump	This indicates a power transition is occurring.	"Server power status is Power Dump" on page 197
Power Termination	The power of the server is starting to transition from on to off.	"Server power status is Power Termination" on page 198
Power Unknown	The state of the server's power is unknown.	"Server power status is Power Unknown" on page 198

Server power status is Power Off

The server power status of Power Off indicates that the power is not applied to a server.

If the power should not be applied to the server, this is good status. If power should be applied to the server, perform the following:

1. Record the server's MTMS from the End-point Topology window in theIBM Network Manager.
2. Go to the HMC that manages the server.
3. Power on the server.
4. If the server will not remain powered on, open Service Focal Point on the HMC that manages the server.
5. Look for a serviceable event that relates to power for that server. Perform the prescribed service for any serviceable event related to that server.
6. If there is no serviceable event that relates to power on that server, and the server is in a 24-inch frame with power supplied by a BPA:
   a. Verify that power is supplied to the BPA and from the BPA to the server.
   b. Perform service on any BPA related serviceable events reported to SFP.
7. If the server is still powered-off, refer to the server documentation for power isolation procedures.

This procedure ends here.

Server power status is Power Transition

The server power status of Power Transition indicates that the power state is changing for server.

Wait a few minutes and check the power status again. The Last Auto-update, displayed in the bottom left of the window, advances in time. If the server power status of Power Transition is persistent, perform the following procedure:

1. Record the server's MTMS from the IBM Network Manager's End-point Topology window.
2. Go to the HMC that manages the server.
3. Power on the server.
4. If the server will not remain powered on, open SFP on the HMC that manages the server.
5. Look for a serviceable event that relates to power for that server. Perform the prescribed service for any serviceable event related to that server.
6. If there is no serviceable event that relates to power on that server, and the server is in a 24-inch frame with power supplied by a BPA:
   a. Verify that power is supplied to the BPA and from the BPA to the server.
   b. Perform service on any BPA related serviceable events reported to SFP.
7. If the server is still powered-off, refer to the server documentation for power isolation procedures.

This procedure ends here.

Server power status is Power Standby

The server power status of Power Standby indicates that the power is at standby. If the power should not be applied to the server, this is good status.

If power should be applied to the server, perform the following:

1. Record the server's MTMS from the IBM Network Manager's End-point Topology window.
2. Go to the HMC that manages the server.
3. Power on the server.
4. If the server will not remain powered on, open SFP on the HMC that manages the server.
5. Look for a serviceable event that relates to power for that server. Perform the prescribed service for any serviceable event related to that server.
6. If there is no serviceable event that relates to power on that server, and the server is in a 24-inch frame with power supplied by a BPA:

a. Verify that power is supplied to the BPA and from the BPA to the server.
b. Perform service on any BPA related serviceable events reported to SFP.

7. If the server is still powered-off, refer to the server documentation for power isolation techniques.

This procedure ends here.

Server power status is Power IPL Transition

The server power status of Power IPL Transition indicates that the power to a server is in a state that has caused the server to begin to IPL.

If the power should not be applied to the server, this is good status. Wait several minutes and check the status again. You should see the Last Auto-update in the bottom left of the window, advance in time. If the Power Status of Power IPL Transition persists, perform the following:

1. Record the server's MTMS from the IBM Network Manager's End-point Topology window.
2. Go to the HMC that manages the server.
3. Power on the server.
4. If the server will not remain powered on, open SFP on the HMC that manages the server.
5. Look for a serviceable event that relates to power for that server. Perform the prescribed service for any serviceable event related to that server.
6. If there is no serviceable event that relates to power on that server, and the server is in a 24-inch frame with power supplied by a BPA:
   a. Verify that power is supplied to the BPA and from the BPA to the server.
   b. Perform service on any BPA related serviceable events reported to SFP.
7. If the server is still powered-off, refer to the server documentation for power isolation procedures.

This procedure ends here.

Server power status is Up

The server power status of Up indicates that power is applied to a server. Generally this is good status.

If the power should not be applied to the server, power off this server. **This procedure ends here.**

Server power status is Power Dump

This status indicates an issue with server power.

If the power should not be applied to the server, this is good status. Wait several minutes and check the status again. You should see the Last Auto-update in the bottom left of the window, advance in time. If the Power Status of Power Dump persists, perform the following:

1. Record the server's MTMS from the IBM Network Manager's End-point Topology window.
2. Go to the HMC that manages the server.
3. Power on the server.
4. If the server will not remain powered on, open SFP on the HMC that manages the server.
5. Look for a serviceable event that relates to power for that server. Perform the prescribed service for any serviceable event related to that server.
6. If there is no serviceable event that relates to power on that server, and the server is in a 24-inch frame with power supplied by a BPA:
   a. Verify that power is supplied to the BPA and from the BPA to the server.
   b. Perform service on any BPA related serviceable events reported to SFP.

7. If the server is still powered-off, refer to the server documentation for power isolation procedures.
8. **This procedure ends here.**

Server power status is Power Termination

The server power status of Power Termination indicates that the server is in transition from Up to Power Off or Power Standby. This should be a temporary state.

If the server was not powered off by a service person, or other user, wait a few minutes and then check the status again. You should see the Last Auto-update in the bottom left of the window, advance in time. If the status is persistently Power Termination, perform the following procedure:

1. Record the server's MTMS from the IBM Network Manager's End-point Topology window.
2. Go to the HMC that manages the server.
3. Power on the server.
4. If the server will not remain powered on, open SFP on the HMC that manages the server.
5. Look for a serviceable event that relates to power for that server. Perform the prescribed service for any serviceable event related to that server.
6. If there is no serviceable event that relates to power on that server, and the server is in a 24-inch frame with power supplied by a BPA:
   a. Verify that power is supplied to the BPA and from the BPA to the server.
   b. Perform service on any BPA related serviceable events reported to SFP.
7. If the server is still powered-off, refer to the server documentation for power isolation techniques.

This procedure ends here.

Server power status is Power Unknown

The server power status of Power Unknown indicates that the IBM Network Manager does not currently know the power status of the server. This could be seen at install or after a repair during the period when the IBM Network Manager is first acquiring information about the status of the server, or this could be the result of a problem with communicating with the server.

Wait a few minutes and check the status again. You should see the Last Auto-update in the bottom left of the window, advance in time. If the switch card power status is persistently Power Unknown, perform the following procedure:

1. Record the server's MTMS from the IBM Network Manager's End-point Topology window.
2. Go to the HMC that manages the server.
3. Power on the server.
4. If the server will not remain powered on, open SFP on the HMC that manages the server.
5. Look for a serviceable event that relates to power for that server. Perform the prescribed service for any serviceable event related to that server.
6. If there is no serviceable event that relates to power on that server, and the server is in a 24-inch frame with power supplied by a BPA:
   a. Verify that power is supplied to the BPA and from the BPA to the server.
   b. Perform service on any BPA related serviceable events reported to SFP.
7. If the server is still powered-off, refer to the server documentation for power isolation procedures.
8. Check the Ethernet connections and network between the server and the HMC that is managing the server, and make any necessary repairs.
9. If the status is still Power Unknown, call your next level of service.

This procedure ends here.

Logical topology window status

This topic discusses the status fields in the Logical Topology window: State and Port State.

For an overview of the Logical topology window, see “Viewing logical topology information in an InfiniBand network” on page 73.

Note: For switches other than 7048-120 or 7048-270, the IBM service representative’s responsibilities for repair and isolation actions stop at the InfiniBand switch. After the IBM service representative has isolated a problem to a switch or a component within a switch chassis, the responsibility of the service representative is to tell the customer to contact their switch vendor’s service organization.

The Logical Topology status window indicates the status of Logical Host Channel Adapters (LHCAs) and Logical Switches (LSW) on a physical Host Channel Adapter (PHCA) and provides the ability for InfiniBand switch partitioning of LHCAs. LHCAs exist only after a PHCA has been assigned to one or more logical partitions. Until the LHCA exists it is not displayed in the Logical Topology view. Because the second port of an LSW is what connects to the LHCA, until an LHCA exists on a PHCA, the Logical Topology View will not display port number 2. This is true for each LSW on any PHCA that does not have an assigned LHCA. Until an LHCA exists on a PHCA, you will see only the PHCA and the first port (which connects to the InfiniBand network) of each LSW; you will not see an LHCA, nor port number 2 of the LSW.

The status fields in the Logical Topology window:

Note: It is possible for the menu option **Selected-Expand/Collapse** and the twistie indicator in the window to get out of synch, where one may act as if the selected object is expanded, while the other may act as if the selected object is collapsed. If this happens, select another object, then reselect the object of interest.

State	Refer to “State in the Logical Topology window”
Port State	Refer to “Port State in a Logical Topology window” on page 200

Note: Servers will be displayed with their MTMS in the location code column, however, HCAs and their ports will have IBM location codes in the location code column.

State in the Logical Topology window

Use the following table to look up the State in the Logical Topology window and cross-reference to the appropriate procedure for that status.

State	Description	Procedure
Not Present	The device is no longer present.	“Adapter status or logical topology state is Not Present” on page 212
Garded	The device has been garded and is not available for use. This is likely caused by an internal adapter error, or bus error.	“Adapter status or logical topology state is Garded” on page 212
Functional	The device is present and functional. This does not imply that the link is available for use; see “Port State in a Logical Topology window” on page 200 for port status.	

Port State in a Logical Topology window

This represents the status of the cable (physical link) for training the link as seen by the port.

State	Description	Procedure
No physical state change	The state has not yet been read by the IBM Network Manager.	"Link training state is No Physical State Change" on page 214
Sleep	This indicates that the port is waiting for a signal from the other side of the cable. This signal will indicate that the link training sequence should begin.	"Link training state is Sleep" on page 214
Polling	This indicates that the port is actively seeking to start the link training sequence with the port on the other side of the cable.	"Link training state is Polling" on page 215
Disabled	The port has been disabled.	"Link training state is Disabled" on page 215
Port configuration training	The port is in the middle of the link training sequence.	"Link training state is Port Configuration Training" on page 215
Linkup	The link has been trained and is available for use.	"Link training state is Linkup" on page 215
Link error recovery	The link has encountered an error condition that requires retraining to recover functionality.	"Link training state is Link Error Recovery" on page 216
Reserved physical state	This state is reserved for future use.	"Link training state is Reserved Physical State" on page 216

Switch topology window status

From the Switch Topology window, there are several ways to view status. There is status in the main window and also through properties for the switch chassis and switch ports.

For an overview of the Switch Topology window, see "Viewing switch topology information in an InfiniBand network" on page 71.

Note: For switches other than 7048-120 or 7048-270, the IBM service representative's responsibilities for repair and isolation actions stop at the InfiniBand switch. After the IBM service representative has isolated a problem to a switch or a component within a switch chassis, the responsibility of the service representative is to tell the customer to contact their switch vendor's service organization.

Use the following table to lookup how to view and interpret the various status fields for switches available through the Switch Topology window

Description	Reference
The main Switch Topology window Port Status	"Port status in the switch topology window" on page 201
The main Switch Topology window Connectivity status	"Connectivity status in switch topology window" on page 202
If you select a switch chassis and choose "Selected-Properties", you can view the chassis LED state by choosing the "System" tab.	"Switch Topology - Switch Properties: System Status" on page 202
If you select a switch port and from the menu choose "Selected-Properties, you can view the port logical status and port physical status in the "System" tab.	"Switch Topology – Port Properties: System Status" on page 203

Description	Reference
If you select a switch chassis and from the menu choose ″Selected-Environmentals″, you can view environmental information about the switch.	“Switch Environmental Status” on page 205

The hierarchy of a port’s status is listed below. The Port Administrate Status must be enabled for the Port Physical state machine to get to the point where the ports on either side of a cable can communicate. Finally, the Port Physical State must be such that the Port Logical State Machine can get to the point where data can actually be passed on the link.

Note: The windows in which you can find the different status are in parentheses, after each type of status.

Port Administrate Status
(Port Administrate in Switch Port Properties window)

Port Physical Status
(Port Connection Status in Switch Port Properties System Status)

Port Logical Status
(Port Status in End-point Topology and Switch Topology windows)

Note: Because the IBM Network Manager does not gather information about PCI HCAs, there will be no PCI HCA port logical status in the Endpoint view. This status only applies to GX bus HCAs.

Note:

1. It is possible for the menu option **Selected-Expand/Collapse** and the twistie indicator in the window to get out of synch, where one may act as if the selected object is expanded, while the other may act as if the selected object is collapsed. If this happens, select another object, then reselect the object of interest.
2. If there are multiple switch ports in a 7048-270 or SFS7008P that are exhibiting trouble, you should first attempt to determine if there is a common failure behind them. Use the procedure found in “Determining faulty fabric controller cards versus faulty LIM cards” on page 153. However, first be sure that there is no common event that may have caused this, like rebooting of many servers in the cluster which are connected to this switch card.

Port status in the switch topology window

State	Description	Procedure
No Port State Change	The state has not yet been read.	“Port status is No Port State Change” on page 213
Down	The link is down.	“Port status is Down” on page 213
Initialized	The link is being initialized.	“Port status is Initialized” on page 213
Armed	The link should be transitioning from initialized to active or down.	“Port status is Armed” on page 213
Active	The link is available for use.	“Port status is Active” on page 213
Active Defer	The link is transitioning from active to down.	“Port status is Active Defer” on page 214
Reserved Port State	This is not a valid state at this time.	“Port status is Reserved Port State” on page 214

Connectivity status in switch topology window

The connectivity status indicates whether IBM Network Manager can communicate with a switch. Bad status can indicate a problem with the service subsystem, or that a switch is powered off. The following table describes the possible connectivity status values:

State	Header	Header
Responsive	IBM Network Manager can communicate with the switch	This is a normal state
Unresponsive	IBM Network Manager cannot communicate with the switch	Use the Isolation Procedure IBNSSLC; see “IBNSSLC” on page 116.

Switch Topology - Switch Properties: System Status

Use this procedure to lookup the state of the chassis LED status in the switch topology window.

This is accessed through the View Switch Topology window, selecting a chassis and then choosing to view the properties and going to the System Status tab. Use the following table to lookup the state for the Chassis LED Status in the switch topology window and cross-reference to the appropriate procedure for that status. The Chassis LED is used by the network manager to provide an LED indentification for a switch. It actually puts the LED into a test mode so that it can override the switch software control of the LED.

State	Description	Procedure
LED on	The IBM Network Manager has set the chassis LED to identify the switch. It will be a flashing yellow.	“Chassis LED status is On”
LED off	The IBM Network Manager is no longer controlling the chassis LED for identification purposes. It will first return to the original state to which the switch software has set it.	“Chassis LED status is Off” on page 203

Chassis LED status is On

The IBM Network Manager is controlling the chassis LED for identification purposes. The LED will be flashing yellow. If you want the network manager to relinquish control over to the switch software, perform the following procedure:

1. In the Switch Topology window, select the switch chassis.
2. From the menu, choose Selected-Identify.
3. Check the Chassis LED status again.
4. If the LED status is still “LED On”, do the following:
   - Check the Ethernet network connections and ports for the service network and verify that they are all operational.
5. **This procedure ends here.**

Chassis LED status is Off

The IBM Network Manager is not controlling the Chassis LED for identification purposes. The LED will be in whatever state the switch software has set. If you want the IBM Network Manager to take control of the LED and set it to a flashing yellow, perform the following procedure:

1. In the Switch Topology window, select the switch chassis.
2. From the menu, choose Selected-Identify.
3. Check the Chassis LED status again.
4. If the LED status is still "LED Off", do the following:
   - Check the Ethernet network connections and ports for the service network and verify that they are all operational.
5. **This procedure ends here.**

Switch Topology – Port Properties: System Status

This topic discusses the switch topology in the Logical Topology Window: Port State.

The Switch Topology – Port Properties window's System tab is accessed through View Switch Topology and then selecting a port and choosing to view the properties.

There are two properties that reflect status: Port Administrate and Port Connection Status.

Description	Reference
The port administrate status indicates the current state that the Subnet Manager has set for the port. The port administrate references the Port Administrate Status, Port Logical Status, and Port Physical Status	"Switch Port System Status Properties: Port Administrate"
The port connection status indicates the physical link state for training the link as seen by the port.	"Switch Port Properties System Status: Port Connection Status" on page 205

Switch Port System Status Properties: Port Administrate

The port administrate status indicates the current state that the Subnet Manager has set for the port.

State	Description	Procedure
No Port Admin State Change	Need more info	"Port Administrate Status is No Port Admin State Change"
Admin Up	This indicates that the Subnet Manager has set the port in a state that will allow it to configure physically and logically.	"Port Administrate Status is Admin Up" on page 204
Admin Down	This indicates that the Subnet Manager has disabled the port. Need to tie this into the physical and logical status procedures.	"Port Administrate Status is Admin Down" on page 204
Admin Testing	The administrate function of this port is under test.	"Port Administrate Status is Admin Testing" on page 204

Port Administrate Status is No Port Admin State Change

The state has not yet been read by the IBM Network Manager.

1. Wait several minutes and check the status again. You must reopen the window to receive an update.

2. If the status is persistently No Port Admin State Change, perform the following procedure:
   a. Verify that the device is powered on.
   b. Check the Ethernet network connectivity to the device over the service network. Verify that all cables are plugged and that all Ethernet switch/router ports appear operational.
3. **This procedure ends here.**

Port Administrate Status is Admin Up

Generally speaking the Port Administrate Status being Admin Up is good status. If you expected this port's administrate status to be Admin Down, perform the following procedure:

1. Select the port.
2. From the menu choose, Selected-Administrate-Disable.
3. Bring up the Port Administrate Status again and verify that the status has changed to Admin Down.
4. **This procedure ends here.**

Port Administrate Status is Admin Down

This indicates that from a Switch management perspective the port is not to be used. If this is not expected, perform the following procedure:

1. Select the port.
2. From the menu choose, Selected-Administrate-Enable.
3. Bring up the Port Administrate Status again and verify that the status has changed to "Admin Up".
4. **This procedure ends here.**

Port Administrate Status is Admin Testing

This indicates that the admin state is under test. This should be a transition state. Perform the following procedure:

1. Wait five (5) minutes and recheck the state. You must reopen the window to get an update.
2. If the state is still "Admin Testing", continue with this procedure. Otherwise, this procedure ends here.
3. Select the port.
4. From the menu choose, Selected-Administrate-Enable.
5. Bring up the Port Administrate Status again and verify that the status has changed to "Admin Up".
6. If the state is still "Admin Testing", you should restart the Subnet manager that controls the subnet of which this port is a part software. See "Checking the subnet manager" on page 140, and "Restarting the subnet manager" on page 142.
7. If after restarting the subnet manager, the port is still in "Admin Testing", you should restart the switch management code on the switch on which this port resides. See "Rebooting the entire switch chassis" on page 145.

   Note: Rebooting the switch management code on a switch can be very disruptive and may cause serviceable events to be generated. You will need to wait about 5 to 10 minutes before any serviceable events created because of the reboot have flowed into SFP. You will note that such events have devices from the rebooted switch chassis in the FRU list. You should close those events, and make a comment that they were opened because of the reboot.
8. **This procedure ends here.**

Switch Port Properties System Status: Port Connection Status

This represents the physical link state for training the link as seen by the port.

State	Description	Procedure
No Physical State Change	The state has not yet been read by the IBM Network Manager.	"Link training state is No Physical State Change" on page 214
Sleep	This indicates that the port is waiting for a signal from the other side of the cable. This signal will indicate that the link training sequence should begin.	"Link training state is Sleep" on page 214
Polling	This indicates that the port is actively seeking to start the link training sequence with the port on the other side of the cable.	"Link training state is Polling" on page 215
Disabled	The port has been disabled.	"Link training state is Disabled" on page 215
Port Configuration Training	The port is in the middle of the link training sequence.	"Link training state is Port Configuration Training" on page 215
Linkup	The link has been trained and is available for use.	"Link training state is Linkup" on page 215
Link Error Recovery	The link has encountered an error condition that requires re-training to recover functionality.	"Link training state is Link Error Recovery" on page 216
Reserved Physical State	This state is reserved for future use.	"Link training state is Reserved Physical State" on page 216

Switch Environmental Status

This topic explains how to view the environmental status for switch devices, such as power supplies, fans, and sensors.

The status for each device is displayed on a separate panel, and there is a separate line for each individual device within the device-type panel. The device-type panel can be accessed through the View Switch Topology window by selecting a switch chassis, and then from the menu choosing **Selected Environmentals**.

Note: The Switch Environmental Status window is also accessed through View Management Properties and choosing the Switch tab. Then you select a switch and click the Environmentals button.

Use "Finding the part in the switch environmental property window" on page 210, to find the device for which the status of interest is being displayed, and use the table below to determine the procedure to follow for that status.

State	Description	Procedure
Normal	Sensors are within normal range.	This is good status.
Warning Temperature	A temperature sensor has reached a warning level.	"Switch environmental status is Warning Temperature" on page 206
Critical Temperature	A temperature sensor has reached a critical level.	"Switch environmental status is Critical Temperature" on page 207

State	Description	Procedure
Voltage Alert	A sensor as reach an over-voltage level.	"Switch environmental status is Voltage Alert" on page 207
Current Alert	A sensor has reached an over-current level.	"Switch environmental status is Current Alert" on page 208
Unknown	The status for the device is unknown.	"Switch environmental status is Unknown" on page 208
Up	The device is up and operational.	"Switch environmental status is Up" on page 209
Down	The device is down and not operational.	"Switch environmental status is Down" on page 209
Failure	The device has experienced a failure.	"Switch environmental status is Failure" on page 210

Switch environmental status is Warning Temperature

The switch environmental status of Warning Temperature indicates that the device is not operating within the normal

1. Record the location code for the switch chassis that you selected and for which you opened the switch environmental property window. Note the portion of the location code that indicates the chassis information: U[*MachineType*].[*Model*].[*SerialNumber*] or U*ProductName.SerialNumber*; where *ProductName* is SFS7000P* or SFS7008P*).
2. Go to SFP and fix any serviceable events against this chassis. If you find a serviceable event related to power or fans, do the following. Otherwise, go to the next step.
   a. After replacing FRUs, according to the isolation procedure referenced in the serviceable event, close and reopen the switch environmental property window.
   b. If the temperature has returned to Normal, **this procedure ends here**. Otherwise, go to the next step.
3. Go to the switch chassis and verify that all the fans are operational by checking for rotation and their LEDs. see "Switch fan tray LEDs" on page 127. If any fan is not operational, do the following procedure. Otherwise, go to the next step.
   a. Replace the FRU containing the faulty fan.
   b. Close and reopen the switch environmental property window.
   c. If the temperature has returned to the Normal state, **this procedure ends here**. Otherwise, go to the next step.
4. Go to the switch chassis indicated by the location code and inspect the area for airflow issues that may impede proper cooling of the switch. If there are any issues:
   a. Remove anything that is impeding proper cooling of the switch.
   b. Close and reopen the switch environmental property window.
   c. If the temperature has returned to the Normal state, **this procedure ends here**. Otherwise, go to the next step.
5. Verify that the site meets the cooling requirements for the switch. See the appropriate hardware guide for the switch model on which you are working. See Table 1 on page 2 for the switch documentation.
6. If the site does not meet cooling requirements, inform the customer what action should be taken to bring the site into compliance. Otherwise, go to the next step.
7. Replace the FRU that contains the sensor indicating Warning Temperature.
8. **This procedure ends here.**

Switch environmental status is Critical Temperature

The switch environmental status of Critical Temperature indicates that the device is not operating within the normal temperature range, and has moved into a critical range. Perform the following procedure:

1. Record the location code for the switch chassis that you selected and for which you opened the switch environmental property window. Note the portion of the location code that indicates the chassis information: U[*MachineType*].[*Model*].[*SerialNumber*] or U*ProductName.SerialNumber*; where *ProductName* is SFS7000P* or SFS7008P*).
2. Go to SFP and fix any serviceable events against this chassis. If you find a serviceable event related to power or fans, do the following. Otherwise, go to the next step.
   a. After replacing FRUs, according to the isolation procedure referenced in the serviceable event, close and reopen the switch environmental property window.
   b. If the temperature has returned to Normal state, **this procedure ends here**. Otherwise, go to the next step.
3. Go to the switch chassis and verify that all the fans are operational by checking for rotation and their LEDs. see "Switch fan tray LEDs" on page 127. If any fan is not operational, do the following procedure. Otherwise, go to the next step.
   a. Replace the FRU containing the faulty fan.
   b. Close and reopen the switch environmental property window.
   c. If the temperature has returned to the Normal state, **this procedure ends here**. Otherwise, go to the next step.
4. Go to the switch chassis indicated by the location code and inspect the area for airflow issues that may impede proper cooling of the switch. If there are any issues:
   a. Remove anything that is impeding proper cooling of the switch.
   b. Close and reopen the switch environmental property window.
   c. If the temperature has returned to the Normal state, **this procedure ends here**. Otherwise, go to the next step.
5. Verify that the site meets the cooling requirements for the switch. See the appropriate hardware guide for the switch model. See Table 1 on page 2 for the switch documentation.
6. If the site does not meet cooling requirements, tell the customer what is needed to bring the site into compliance. **This procedure ends here.** Otherwise, go to the next step.
7. Replace the FRU that contains the sensor indicating Critical Temperature.
8. **This procedure ends here.**

Switch environmental status is Voltage Alert

The switch environmental status of Voltage Alert indicates that the device is not operating within the normal voltage range. Do the following procedure:

1. Record the location code for the switch chassis that you selected and for which you opened the switch environmental property window. Note the portion of the location code that indicates the chassis information: U[*MachineType*].[*Model*].[*SerialNumber*] or U*ProductName.SerialNumber*; where *ProductName* is SFS7000P* or SFS7008P*).
2. Go to SFP and fix any serviceable events against this chassis. If you find a serviceable event related to power or fans, do the following. Otherwise, go to the next step.
   a. After replacing FRUs, according to the isolation procedure referenced in the serviceable event, close and reopen the switch environmental property window.
   b. If the status has returned to Normal state, **this procedure ends here**. Otherwise, go to the next step.
3. Go to the switch chassis and verify that all the power supplies are operational by checking their LEDs. see "7048-270 or SFS7008P switch power supply LEDs" on page 126. If any power supply is not operational, do the following procedure. Otherwise, go to the next step.

a. Replace the FRU that contains the failing power supply.
b. Close and reopen the switch environmental property window.
c. If the status has returned to the Normal state, **this procedure ends here**. Otherwise, go to the next step.

4. Replace the FRU that contains the sensor indicating Voltage Alert, and do the following procedure.
   a. Close and reopen the switch environmental property window.
   b. If the status has returned to the Normal state, **this procedure ends here**. Otherwise, go to the next step.
5. Verify that the site meets the power and cooling requirements for the switch. See the appropriate hardware guide for the switch. See Table 1 on page 2 for the switch documentation.
6. If the site does not meet power and cooling requirements, tell the customer what is needed to bring the site into compliance. **This procedure ends here.**

Switch environmental status is Current Alert

The switch environmental status of Current Alert indicates that the device is not operating within the normal current range.

1. Record the location code for the switch chassis that you selected and for which you opened the switch environmental property window. Note the portion of the location code that indicates the chassis information: U[*MachineType*].[*Model*].[*SerialNumber*] or `U`*`ProductName`*`.`*`SerialNumber`*; where *ProductName* is SFS7000P* or SFS7008P*).
2. Go to SFP and fix any serviceable events against this chassis. If you find a serviceable event related to power or fans, do the following. Otherwise, go to the next step.
   a. After replacing FRUs, according to the isolation procedure referenced in the serviceable event, close and reopen the switch environmental property window.
   b. If the status has returned to Normal state, **this procedure ends here**. Otherwise, go to the next step.
3. Go to the switch chassis and verify that all the power supplies are operational by checking their LEDs. see "7048-270 or SFS7008P switch power supply LEDs" on page 126. If any power supply is not operational, do the following procedure. Otherwise, go to the next step.
   a. Replace the FRU that contains the failing power supply.
   b. Close and reopen the switch environmental property window.
   c. If the status has returned to the Normal state, **this procedure ends here**. Otherwise, go to the next step.
4. Replace the FRU that contains the sensor indicating Voltage Alert, and do the following procedure.
   a. Close and reopen the switch environmental property window.
   b. If the status has returned to the Normal state, **this procedure ends here**. Otherwise, go to the next step.
5. Verify that the site meets the power and cooling requirements for the switch. See the appropriate hardware guide for the switch. See Table 1 on page 2 for the switch documentation.
6. If the site does not meet power and cooling requirements, tell the customer what is needed to bring the site into compliance. **This procedure ends here.**

Switch environmental status is Unknown

This indicates that the environmental status for a device is not accessible.

1. Wait a few minutes and check the status again. You must reopen the window to receive the update.
2. If the switch environmental status for the devices is persistently Unknown, perform the following procedure:

a. Record the location code for the switch chassis that you selected and for which you opened the switch environmental property window. Note the portion of the location code that indicates the chassis information: U[*MachineType*].[*Model*].[*SerialNumber*] or U*ProductName.SerialNumber*; where *ProductName* is SFS7000P* or SFS7008P*).
b. Go to SFP on the HMC that is running the IBM Network Manager
c. Look for any serviceable event that was created within that chassis. Look for FRUs with a location code that begins with the location for the chassis to do this. All devices in the same chassis will have a location code that begins with the location for the chassis and then continue with card specific information.
d. If there is a serviceable event against a card in the chassis, perform the prescribed service.
   1) After performing the prescribed service, recheck the status.
   2) If the status is still Unknown, then continue with the next step. Otherwise, this procedure ends here.
e. Check the Ethernet connections and network between the switch and the HMC that is running the IBM Network Manager, and make any necessary repairs.
f. If the status is still Unknown, call your next level of service.

3. **This procedure ends here.**

Switch environmental status is Up

The switch environmental status of Up indicates good status.

Switch environmental status is Down

The switch environmental status of Down indicates that the device is not operational. Perform the following procedure:

1. Record the location code for the switch chassis that you selected and for which you opened the switch environmental property window. Note the portion of the location code that indicates the chassis information: U[*MachineType*].[*Model*].[*SerialNumber*] or U*ProductName.SerialNumber*; where *ProductName* is SFS7000P* or SFS7008P*).
2. Go to SFP on the HMC that is running the IBM Network Manager
3. Look for any serviceable event that was created within that chassis. Look for FRUs with a location code that begins with the location for the chassis to do this. All devices in the same chassis will have a location code that begins with the location for the chassis and then continue with card specific information.
4. Do one of the following actions:
   - If there is a serviceable event against a card in the chassis, perform the prescribed service.
     a. After performing the prescribed service, recheck the status.
     b. If the status is still Down, then continue with the next step. Otherwise, this procedure ends here.
   - If there is no serviceable event reported to SFP, do one of the following:

     Note: In the following steps, U[*unit location*] can be either U[*MachineType*].[*Model*].[*SerialNumber*] or U*ProductName.SerialNumber*; where *ProductName* is SFS7000P* or SFS7008P*,

     – If the status is in the Power Supplies section, replace the card located in U[*unitlocation*]P1-E*y*; where *y* is the ID.
     – If the status is in the Fans section, do one of the following:
       - If the switch is a 7048-120 or SFS7000P, replace the card located in U[*unitlocation*]P1-E*x*; where *y* is 1 for IDs 1 and 2, and *y* is 2 for IDs 3 and 4.
       - If the switch is a 7048-270 or SFS7008P, replace the card located in U[*unitlocation*]P1-A*y*; where *y* is 1 for IDs 1 and 2, and *y* is 2 for IDs 3 and 4.

Note: The 7048-120 or SFS7000P fans are on the power supply cards. The 7048-270 or SFS7008P has fan trays separate from the power supplies.

– If the status is in the Sensors section, replace the card located in U[*unitlocation*]P1-C*y*.

5. Replace the part.
6. If the status is still Down, call your next level of service.
7. **This procedure ends here.**

Switch environmental status is Failure

The switch environmental status of Failure indicates that the device has experienced a failure. Perform the following procedure:

1. Record the location code for the switch chassis that you selected and for which you opened the switch environmental property window. Note the portion of the location code that indicates the chassis information: U[*MachineType*].[*Model*].[*SerialNumber*] or U*ProductName.SerialNumber*; where *ProductName* is SFS7000P* or SFS7008P*).
2. Go to SFP on the HMC that is running the IBM Network Manager
3. Look for any serviceable event that was created within that chassis. Look for FRUs with a location code that begins with the location for the chassis to do this. All devices in the same chassis will have a location code that begins with the location for the chassis and then continue with card specific information.
4. Do one of the following actions:
   - If there is a serviceable event against a card in the chassis, perform the prescribed service.
     a. After performing the prescribed service, recheck the status.
     b. If the status is still Failure, then continue with the next step. Otherwise, this procedure ends here.
   - If there is no serviceable event reported to SFP, do one of the following:

     Note: In the following steps, U[*unit location*] can be either U[*MachineType*].[*Model*].[*SerialNumber*] or U*ProductName.SerialNumber*; where *ProductName* is SFS7000P* or SFS7008P*,

     – If the status is in the Power Supplies section, replace the card located in U[*unitlocation*]P1-E*y*; where *y* is the ID.
     – If the status is in the Fans section, do one of the following:
       - If the switch is a 7048-120 or SFS7000P, replace the card located in U[*unitlocation*]P1-E*y*; where *y* is 1 for IDs 1 and 2, and *y* is 2 for IDs 3 and 4.
       - If the switch is a 7048-270 or SFS7008P, replace the card located in U[*unitlocation*]P1-A*y*; where *y* is 1 for IDs 1 and 2, and *y* is 2 for IDs 3 and 4.

       Note: The 7048-120 or SFS7000P fans are on the power supply cards. The 7048-270 or SFS7008P has fan trays separate from the power supplies.

     – If the status is in the Sensors section, replace the card located in U[*unitlocation*]P1-C*y*.
5. Replace the part.
6. If the status is still Failure, call your next level of support.
7. **This procedure ends here.**

Finding the part in the switch environmental property window

The switch environmental property window has status for switch power supplies, fans and sensors. If you have bad environmental status recorded against a device and can't find a serviceable event against the device, you may be directed to replace it. The following table has a column for 7048-120 or SFS7000P locations and one for 7048-270 or SFS7008P locations. Each location is described with both prose and a location code.

Device	7048-120 or SFS7000P Location	7048-270 or SFS7008P Location
Power Supply 1	U7048.120.[sn]-P1-E1 USFS7000P*.[sn]-P1-E1 Power/Fan module on the left	U7048.270.[sn]-P1-E1 USFS7008P*.[sn]-P1-E1 Power module on the left
Power Supply 2	U7048.120.[sn]-P1-E2 USFS7000P*.[sn]-P1-E2 Power/Fan module on the right	U7048.270.[sn]-P1-E2 USFS7008P*.[sn]-P1-E1 Power module on the right
Fan 1	U7048.120.[sn]-P1-E1 USFS7000P*.[sn]-P1-E1 Part of power supply 1.	U7048.270.[sn]-P1-A1 USFS7008P*.[sn]-P1-A1 Fan Tray on the Left
Fan 2	U7048.120.[sn]-P1-E2 USFS7000P*.[sn]-P1-E2 Part of power supply 2.	U7048.270.[sn]-P1-A2 USFS7008P*.[sn]-P1-A1 Fan Tray on the Right
Sensor X	U7048.120.[sn]-P1-C1 USFS7000P*.[sn]-P1-C1 The temperature sensor on the switch card.	U7048.270.[sn]-P1-C*x* USFS7008P*.[sn]-P1-C*x* where *x* corresponds to the slot ID displayed The temperature sensor on a switch card.

Common Status Procedures

Contains procedures referring to values for status fields that are shared across multiple IBM Network Manager GUI windows.

These procedures refer to the values for status fields that are shared across multiple IBM Network Manager GUI windows.

The following table contains the subtopics and procedures found in Common Status Procedures.

“Adapter status or logical topology state is Functional” on page 212	“Port status is Reserved Port State” on page 214
“Adapter status or logical topology state is Not Present” on page 212	“Procedures for link training states as seen by a port” on page 214
“Adapter status or logical topology state is Garded” on page 212	“Link training state is No Physical State Change” on page 214
“Procedures for link state as seen by a port” on page 212	“Link training state is Sleep” on page 214
“Port status is No Port State Change” on page 213	“Link training state is Polling” on page 215
“Port status is Down” on page 213	“Link training state is Disabled” on page 215
“Port status is Initialized” on page 213	“Link training state is Port Configuration Training” on page 215
“Port status is Armed” on page 213	“Link training state is Linkup” on page 215
“Port status is Active” on page 213	“Link training state is Link Error Recovery” on page 216
“Port status is Active Defer” on page 214	“Link training state is Reserved Physical State” on page 216

Adapter status or logical topology state is Functional

This indicates good status.

Adapter status or logical topology state is Not Present

If the adapter status is Not Present and you expect it to be functional, perform the following actions:

1. Record the location code for the adapter.
2. Go to the controlling HMC for the server in which the adapter is populated.
3. Check for a serviceable event with that location code in the FRU list.
   a. If there is a serviceable event logged against that location code, take the appropriate action for that serviceable event.
   b. If there is no serviceable event logged against that location code, go to the next step.
4. Using the location code provided go to the adapter and check the LEDs.
   a. If there is no adapter present at that location, you will need to install one for the status to go to Functional. Before installing an adapter, you should first determine if there is a valid reason for the adapter not being present. Valid reasons include: (1) waiting for a replacement FRU (2) the customer has decided to move the adapter.
   b. If there is an adapter present at the location, and the server is powered on, go to the appropriate server or adapter documentation "InfiniBand switch reference information" on page 3 and determine if the LEDs on the adapter indicate a problem. If they do, fix the problem according to the server or adapter documentation instructions. Depending on the status of the adapter, it is possible that when you remove the cable, a serviceable event will be reported by the switch through to SFP. Therefore, be sure to record the time at which you removed the cable from the adapter port, and also take into account the differences between the time source you use for recording the cable removal and the HMC on which the IBM Network Manager is running; see "Understanding timestamp differences" on page 145.
   c. If there is an adapter present and the LEDs do not indicate a problem, you should query the adapter status using the O/S to see if it is seen by the O/S and available for use by the O/S. Use the appropriate server or adapter manual "InfiniBand switch reference information" on page 3 to determine how to resolve any issue with adapter availability to the O/S.
   d. **This procedure ends here.**

Adapter status or logical topology state is Garded

If the adapter status is Guarded and you expect it to be functional, perform the following actions:

1. Record the location code for the adapter.
2. Go to SFP on the controlling HMC for the server in which the adapter is populated.
3. Look for a serviceable event logged with the location code for the adapter in the FRU list.
   a. If there is a serviceable event logged with the adapter in the FRU list, service the event in the prescribed manner.
   b. If there is not a serviceable event logged with the adapter in the FRU list, call your next level of support.
4. **This procedure ends here.**

Procedures for link state as seen by a port

The link states seen by a port for which you will find procedures in the following sub-sections are logical link states, as opposed to the more physical link training states described in another section. The following table describes in which windows you may see these states.

Table 23.

Window	Field
End-point Topology	Port Status
Switch Topology	Port Status

Port status is No Port State Change

The state has not yet been read by the IBM Network Manager.

1. Wait several minutes and check the status again. You should see the Last Auto-update in the bottom left of the window advance in time.
2. If the status is persistently No Physical State Change, perform the following procedure:
   a. Verify that the device is powered on.
   b. Check the Ethernet network connectivity to the device over the service network. Verify that all cables are plugged and that all Ethernet switch/router ports appear operational.
3. **This procedure ends here.**

Port status is Down

This indicates that a port has gone down.

Perform the procedure in section "Isolating a problem with a port or link" on page 136.

Note: If the switch port has a 4x connector from an octopus cable attached to it, you should refer to the following section to determine if the port should be the Active Port: "Verifying static-12x mode connectivity" on page 67. If you want to verify if the port is configured for static-12x operation, see "Verifying Static12x or 4x configuration to a port" on page 153.

Port status is Initialized

This should be an intermediate state when a port is first activated.

1. The link is in the process of initializing, wait a few minutes and check the status again. You should see the Last Auto-update in the bottom left of the window advance in time.
2. If the status is persistently Initialized, perform the procedure in "Isolating a problem with a port or link" on page 136.
3. If the problem persists, call your next level of support.
4. **This procedure ends here.**

Port status is Armed

This should be an intermediate state when a port is first activated.

1. The link is in the process of initializing, wait a few minutes and check the status again. You should see the Last Auto-update in the bottom left of the window advance in time.
2. If the status is persistently Armed, perform the procedure in "Isolating a problem with a port or link" on page 136.
3. If the problem persists, call your next level of support.
4. **This procedure ends here.**

Port status is Active

This is good status for a port.

Note: If the switch port has a 4x connector from an octopus cable attached to it, you should refer to the following section to determine if the port should be the Active Port: “Verifying static-12x mode connectivity” on page 67. If you want to verify if the port is configured for static-12x operation, see “Verifying Static12x or 4x configuration to a port” on page 153.

Port status is Active Defer

This should indicates that a state is going from Active to Down.

1. Wait a few minutes and check the status again. You should see the Last Auto-update in the bottom left of the window advance in time.
2. If the status is persistently Active Defer, perform the procedure in “Isolating a problem with a port or link” on page 136.
3. If the problem persists, call your next level of support.
4. **This procedure ends here.**

Port status is Reserved Port State

This state should not be entered. Call your next level of support.

Procedures for link training states as seen by a port

The following subsections describe what to do when you encounter these link training states. This is displayed in the following fields:

Window	Field
Logical Topology	Port State
Switch Topology – Port Properties - System tab	Port Connection Status

Link training state is No Physical State Change

The state has not yet been read by the IBM Network Manager.

1. Wait several minutes and check the status again. You must reopen the window to receive an update.
2. If the status is persistently No Physical State Change, perform the following procedure:
   a. Verify that the device is powered on.
   b. Check the Ethernet network connectivity to the device over the service network. Verify that all cables are plugged and that all Ethernet switch or router ports appear operational.
3. **This procedure ends here.**

Link training state is Sleep

The port is in the sleep state where it is waiting for a signal from a port on the other side of the link before starting the link training sequence. You may see this after an error condition, or on install, or after a repair action, or if a link has otherwise been taking out of service and put back into service.

1. If the link should be in service, wait a few minutes and check the status again. You must reopen the window to receive an update.
2. If the status is persistently Sleep, check for the existence of a cable on the link:
   a. Record the location code for the device’s port and open the properties and record the location code for the port on the other side of the link.

      If the properties do not indicate the remote port to which this port should be connected, then either get the information from a label on the cable, or from any available cable planning documentation.

b. Look for a cable connected to both ports, and verify that it is connected between the two ports. The most reliable way to do this is to make sure that all cable ends have been connected to a port.

3. If a cable exists on the link, perform the procedure in "Isolating a problem with a port or link" on page 136.
4. If the problem persists, call your next level of support.
5. This procedure ends here.

Link training state is Polling

This indicates that the port is polling for the existence of another port on the link. You may see this after an error condition, or on install, or after a repair action, or if a link has otherwise been taking out of service and put back into service.

1. If the link should be in service, wait a few minutes and check the status again. You must reopen the window to receive an update.
2. If the status is persistently "Polling", check for the existence of a cable on the link:
   a. Record the location code for the device's port and open the properties and record the location code for the port on the other side of the link.

      If the properties do not indicate the remote port to which this port should be connected, then either get the information from a label on the cable, or from any available cable planning documentation.

   b. Look for a cable connected to both ports, and verify that it is connected between the two ports. The most reliable way to do this is to make sure that all cable ends have been connected to a port.
3. If a cable exists on the link, perform the procedure in "Isolating a problem with a port or link" on page 136.
4. If the problem persists, call your next level of support.
5. This procedure ends here.

Link training state is Disabled

This indicates that the port has been disabled and is not usable. You may see this under after an error condition, or on install, or after a repair action, or if a link has otherwise been taking out of service and put back into service.

1. Wait a few minutes and check the status again. You must reopen the window to receive an update.
2. If the status is persistently Disabled, perform the procedure in "Isolating a problem with a port or link" on page 136.
3. If the problem persists, call your next level of support.
4. **This procedure ends here.**

Link training state is Port Configuration Training

This indicates that the port is in the process of training the link. You may see this on install, or after a repair action, or if a link has otherwise been taking out of service and put back into service.

1. Wait a few minutes and check the status again. You must reopen the window to receive an update.
2. If the status is persistently Port Configuration Training, perform the procedure in "Isolating a problem with a port or link" on page 136.
3. If the problem persists, call your next level of support.
4. **This procedure ends here.**

Link training state is Linkup

This is a good state.

Link training state is Link Error Recovery

This indicates that the link has encountered an error affecting the training of the link. This should transition to another state as part of the recovery.

1. Wait a few minutes and check the status again. You must reopen the window to receive an update.
2. If the status is persistently Link Error Recovery, perform the procedure in "Isolating a problem with a port or link" on page 136.
3. If the problem persists, call your next level of support.
4. **This procedure ends here.**

Link training state is Reserved Physical State

This state should not be seen. Call your next level of support.

InfiniBand component location codes

This topic identifies and defines the location codes used by the InfiniBand switch. The FRU parts have their own location codes that are described in this section.

InfiniBand parts have their own location codes. The location codes are defined as follows:

Valid switch FRU location codes

A valid switch logic FRU has one of the following formats:

U*machine type*.*model*.*serial number*-P*x*-C*y*-T*z*
U*product name*.*serial number*-P*x*-C*y*-T*z*

where:
- *x* is equal to the planar number of a backplane (always present)
- *y* is equal to the card slot number of the card (not present for a backplane)
- *z* is equal to the cable connector number on the card (only valid for ports or cables)

A valid switch power FRU has one of the following formats:

U*machine type*.*model*.*serial number*-P*x*-E*power supply identifier*
U*product name*.*serial number*-P*x*-E*power supply identifier*

where:
- *x* is equal to the planar number of a backplane (always present)

A valid switch fan FRU has one of the following formats. If the fan is on the same FRU as the power supply, then it takes on the format of a power supply FRU.

U*machine type*.*model*.*serial number*-P*x*-A*fan identifier*
U*product name*.*serial number*-P*x*-A*fan identifier*

where:
- *x* is equal to the planar number of a backplane (always present)

Valid adapter FRU location codes

A valid adapter FRU location code has the following format:

U*machine type or feature code*.001.*serial number*-P*x*-C*y*-T*z*

where:
- *x* = planar number of a backplane (always present)
- *y* = card slot number of the card (always present)
- *z* = cable connector number on the card (only valid for ports or cables)

Place holders in partial location codes

A location code may have placeholders for certain location fields, because the precise number for that part of the location code may not be known. A pound-sign (#) is used as a place holder.

If you have the switch machine type, model and serial number or the system unit feature code and serial number you can find the enclosure in which the device resides. To determine the precise card and port locations, use the Reference Code Extension information from the FRU list in Service Focal Point. See How to find a Switch FRU using the Reference Code Extension

© Copyright IBM Corp. 2004, 2007

Procedures for finding FRUs

Help to find the correct failing FRU.

In order to find a FRU you must first understand the information that you have available to identify it. The first key is the FRU location code, which is obtained from the FRU list in Service Focal Point.

If you do not have a valid switch or adapter location code, you will have to perform a few more steps to get to the point where you understand the location of the device.

Use the following table to help find which procedure to use:

Symptom	Desired FRU	
	Switch FRU	**Adapter FRU**
Valid Location Code for Desired FRU	"How to find a switch FRU with a valid location code" on page 222.	"How to find an adapter FRU with a valid location code" on page 222
Valid Location for Device to which FRU is connected	"How to find a switch FRU using another device with a valid location code" on page 223	"How to find an adapter FRU using another device with a valid location code" on page 223
Reference Code Extension	"How to find a switch FRU using a reference code extension" on page 224	"How to find an adapter FRU using a reference code extension" on page 176
Using the Error Description	"How to determine a switch FRU using the error description" on page 227	"How to determine an adapter FRU using the error description" on page 227

Location codes used by the IBM Network Manager

Location codes used by the IBM Network Manager

The following subsections give the valid switch FRU location code formats and a physical description of where the parts are located.

Valid Switch FRU Location Codes

A valid switch logic FRU has one of the following formats:

U*machine_type*.*model*.*serial_number*-P*x*-C*y*-T*z*

Machine type, model, and serial number are based on the MTMS of the chassis for a 7048-120 or 7048-270.
x = planar number of a backplane (always present)
y = card slot number of the card (not present for a backplane)
z = cable connector number on the card (only valid for ports, or cables)

U*product_name*.*serial_number*-P*x*-C*y*-T*z*

Product name and serial number are based on the model of the switch, and are used for Cisco switches.
x = planar number of a backplane (always present)
y = card slot number of the card (not present for a backplane)
z = cable connector number on the card (only valid for ports, or cables)

A valid switch power FRU has one of the following formats:

U*machine_type*.*model*.*serial_number*-Px-E*power_supply_id*]

Machine type, model, and serial number are based on the MTMS of the chassis.
x = planar number of a backplane (always present)

A valid switch fan FRU has the following format. If the fan is on the same FRU as the power, then it takes on the format of a power FRU.

U*machine_type.model.serial_number*-P*x*-A*fan_id*]

Machine type, model, and serial number are based on the MTMS of the chassis.

x = planar number of a backplane (always present)

Switch FRU Locations for the Topspin 120 Server Switch (7048-120) or a Cisco 7000P Server Switch (SFS7000P)

There are only three FRUs located in a 7048-120 or SFS7000P. They are the chassis and the two power/fan modules. However, there are an additional 14 FRU location codes used for identifying connectors.

The following figure shows the front of the 7048-120 or SFS7000P and the FRUs that are accessible from the front.

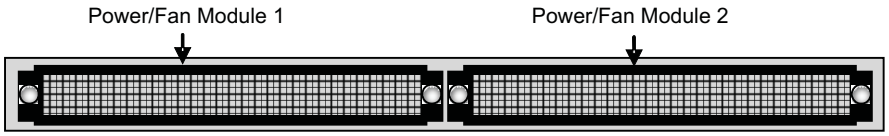

Front View Reference	Location Code
Power/Fan Module 1	U7048.120.*serial_number*-P1-E1 USFS7000P*.*serial_number*-P1-E1
Power/Fan Module 2	U7048.120.*serial_number*-P1-E2 USFS7000P*.*serial_number*-P1-E1

Using the Rear-view figure of the 7048-120 or SFS7000P and the table below it, you can find the FRUs that are accessible from the rear of the 7048-120 or SFS7000P.

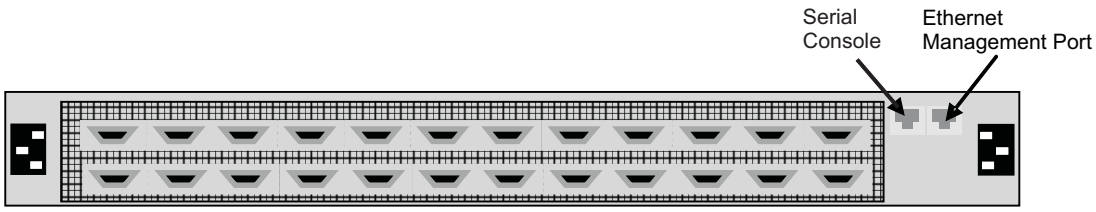

Front View Reference	Location Code
Chassis	U7048.120.*serial_number*-P1 USFS7000P*.*serial_number*-P1
Switch Card (FRU = switch chassis)	U7048.120.*serial_number*-P1-C1 USFS7000P*.*serial_number*-P1-C1
Serial Console Port (FRU = switch chassis)	U7048.120.*serial_number*-P1-T1 USFS7000P*.*serial_number*-P1-T1
Ethernet Management Port (FRU = switch chassis)	U7048.120.*serial_number*-P1-T2 USFS7000P*.*serial_number*-P1-T2
Switch Card Connector, where x is 1 to 24, and connector numbering is sequential from left to right. This is used to indicate the cable end connected to this connector.	U7048.120.*serial_number*-P1-C1-T*x* USFS7000P*.*serial_number*-P1-C1-T*x*

Switch FRU Locations for the 7048-270 or SFS7008P

There are ten FRUs located in a 7048-270 or SFS7008P. They are the chassis and the two power/fan modules.

The following figure shows the front of the 7048-270 or SFS7008P and the FRUs that are accessible from the front.

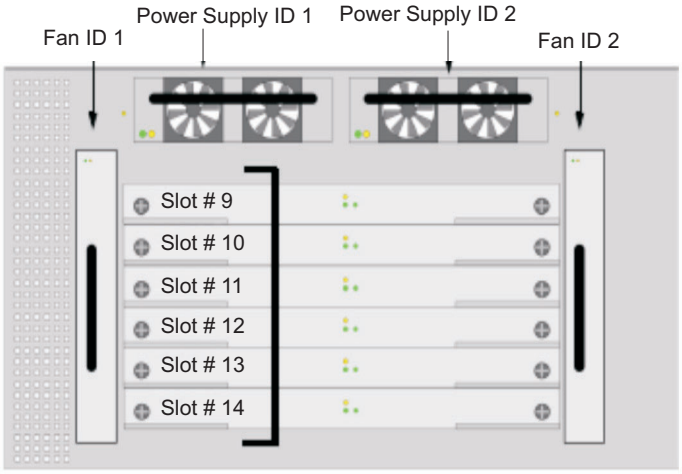

Front View Reference	Location Code Location Code
Fabric Controller Module Node1 in Slot 9	U7048.270.*serial_number*-P1-C9 USFS7008P*.*serial_number*-P1-C9
Fabric Controller Module Node2 in Slot 10	U7048.270.*serial_number*-P1-C10 USFS7008P*.*serial_number*-P1-C10
Fabric Controller Module Core1 Slot 11	U7048.270.*serial_number*-P1-C11 USFS7008P*.*serial_number*-P1-C11
Fabric Controller Module Core2 Slot 12	U7048.270.*serial_number*-P1-C12 USFS7008P*.*serial_number*-P1-C12
Fabric Controller Module Node3 Slot 13	U7048.270.*serial_number*-P1-C13 USFS7008P*.*serial_number*-P1-C13
Fabric Controller Module Node4 Slot 14	U7048.270.*serial_number*-P1-C14 USFS7008P*.*serial_number*-P1-C14
Power Supply ID 1	U7048.270.*serial_number*-P1-E1 USFS7008P*.*serial_number*-P1-E1
Power Supply ID 2	U7048.270.*serial_number*-P1-E2 USFS7008P*.*serial_number*-P1-E2
Fan ID 1	U7048.270.*serial_number*-P1-A1 USFS7008P*.*serial_number*-P1-A1
Fan ID 2	U7048.270.*serial_number*-P1-A2 USFS7008P*.*serial_number*-P1-A2

The following figure shows the rear of the 7048-270 or SFS7008P and the FRUs that are accessible from the rear.

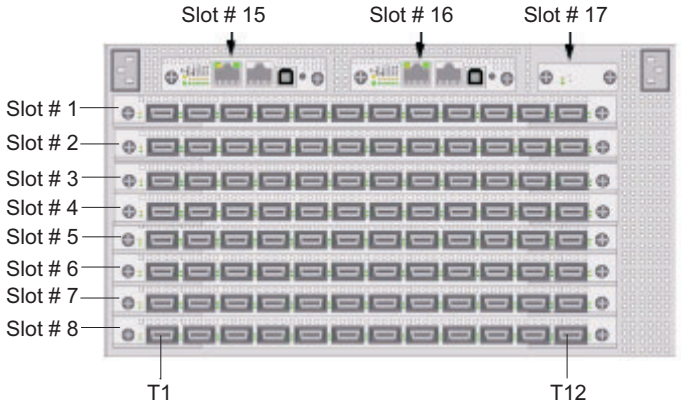

Rear View Reference	Location Code
Connectors on the Line Interface Modules (LIM), where x represents slots 1 through 8, and y represents connectors 1 through 12 on a LIM. The connectors are numbered sequentially from left to right. This is used to indicate the cable end connected to this connector.	U7048.270.*serial_number*-P1-C*x*-T*y* USFS7008P*.*serial_number*-P1-C*x*-T*y*
Line Interface Module (LIM) in Slot 1	U7048.270.*serial_number*-P1-C1 USFS7008P*.*serial_number*-P1-C1
Line Interface Module (LIM) in Slot 2	U7048.270.*serial_number*-P1-C2 USFS7008P*.*serial_number*-P1-C2
Line Interface Module (LIM) in Slot 3	U7048.270.*serial_number*-P1-C3 USFS7008P*.*serial_number*-P1-C3
Line Interface Module (LIM) in Slot 4	U7048.270.*serial_number*-P1-C4 USFS7008P*.*serial_number*-P1-C4
Line Interface Module (LIM) in Slot 5	U7048.270.*serial_number*-P1-C5 USFS7008P*.*serial_number*-P1-C5
Line Interface Module (LIM) in Slot 6	U7048.270.*serial_number*-P1-C6 USFS7008P*.*serial_number*-P1-C6
Line Interface Module (LIM) in Slot 7	U7048.270.*serial_number*-P1-C7 USFS7008P*.*serial_number*-P1-C7
Line Interface Module (LIM) in Slot 8	U7048.270.*serial_number*-P1-C8 USFS7008P*.*serial_number*-P1-C8
Management I/O Card in Slot 15	U7048.270.*serial_number*-P1-C15 USFS7008P*.*serial_number*-P1-C15
Management I/O Card in Slot 16	U7048.270.*serial_number*-P1-C16 USFS7008P*.*serial_number*-P1-C16
Chassis ID Card in Slot 17	U7048.270.*serial_number*-P1-C17 USFS7008P*.*serial_number*-P1-C17

Valid Adapter FRU Location Codes

A valid adapter FRU location code has the following format:

U*server_feature_code*.001.*server_serial_number*-P*x*-C*y*-T*z*

Placeholders in Partial Location Codes

A location code may have placeholders for certain location fields, because the precise number for that part of the location code may not be known. A pound-sign (#) is used as a place holder.

If you have the switch machine type, model and serial number or the CEC feature code and serial number you can find the cage in which the device resides. Beyond that you will need to use the Reference Code Extension information from the FRU list to determine the precise card and port locations. See “How to find a switch FRU using a reference code extension” on page 224, or “How to find an adapter FRU using a reference code extension” on page 176.

How to find a switch FRU with a valid location code

Find a switch FRU using a valid location code

If a switch FRU has a valid Location Code, you can use that to find the FRUs physical location by understanding the format of the location code and cross-referencing information in the View Switch Topology window of the IBM Network Manager GUI.

1. Find and record Information about the FRU:
   a. Record the location code
   b. Record the machine type, model and serial number (MTMS) from the location code
   c. Open the IBM Network Manager GUI’s View Switch Topology window.
   d. Record the chassis name and frame.
2. 2. Turn on Identify LEDs:
   a. Click on the device that matches the recorded location code in the Location-Code field.
   b. On the menu, choose **Selected-Identify-On**
3. Go to physical location and verify FRU

   If you know the physical location of the switch by the Chassis name, the MTMS, or the frame, go to that switch, and verify that the Identify LED is on. If you don’t know the physical location based on the Chassis name, MTMS, or frame, you will have to walk the floor and look for Identify LEDs that are on and then verify that the switch at which you are looking matches the MTMS for the FRU.

How to find an adapter FRU with a valid location code

Find an adapter FRU using a valid location code

If an adapter FRU has a valid Location Code, you can use that to find the its physical location by understanding the format of the location code and cross-referencing information in the View Switch Topology window of the IBM Network Manager GUI.

1. Find and record Information about the FRU:
   a. Record the location code
   b. Record the server machine type, model and serial number (MTMS) from the location code
   c. Open the IBM Network Manager to its View End-Point.
   d. Record the server name and frame:cage information. If this is a 19-inch rack, the cage information is not needed.
2. Turn on Identify LEDs:
   a. Go to the controlling HMC for the server.
   b. Using Service Focal Point-Service Utilities turn on the Identify LED for the server.
3. Go to physical location and verify FRU:
   - If you know the physical location of the server by the Chassis name, the MTMS, or the frame, go to that server, and verify that the identify LED is on.
   - If you don’t know the physical location based on the chassis name, MTMS, or frame, you will have to walk the floor and look for identify LEDs that are on and then verify that the switch at which you are looking matches the MTMS for the FRU.

How to find a switch FRU using another device with a valid location code

How to find a Switch FRU using another Device with a Valid Location Code

Sometimes there is a problem with the location code of a FRU. In such cases, you may be able to use the location code of a FRU to which the desired FRU is attached. This works under the assumption that none of the other FRUs in the same chassis with the FRU with the invalid location will have valid location codes.

It is possible that the location code for the desired FRU is now available in the IBM Network Manager.

Note: For this procedure to work, the network should not have been re-cabled since the time of the latest occurrence of the serviceable event. It is particular important that the failing link was not re-cabled.

1. Find and record Information about the FRU:
   a. Record the location code for the known device.
   b. If the known device has port information in the location code, proceed to the next step. Otherwise, you will need to go to "How to find a switch FRU using a reference code extension" on page 224.
   c. Proceed only if you have a known device with port information.
   d. If the known device is switch port, perform the following procedure. Otherwise, proceed to the next step.
      1) Record the location code of the known device.
      2) Open the IBM Network Manager to the View Switch Topology window.
      3) Using its location code, find the known device in the window.
      4) Record the Neighbor Location-Code, Neighbor Name, and Neighbor Frame:Cage.
   e. You have determined that the known device is an adapter port, so perform the following procedure to find the server that contains the adapter:
      1) Record the location code of the known device
      2) Open the IBM Network Manager to the View End-Point window.
      3) Using its location code, find the known device in the window.
      4) Record the Neighbor Location-Code, Neighbor Name, and Neighbor Frame:Cage.
2. Turn on Identify LEDs:
   a. Go to the IBM Network ManagerView Switch Topology window.
   b. Click on the device that most closely matches the Neighbor location code, name and frame:cage information gathered previously; this should be the desired switch FRU.
   c. On the menu, choose **Selected-Identify-On**
3. Go to physical location and verify FRU:
   - If you know the physical location of the desired switch by the name, the MTMS, or the frame, go to that switch, and verify that the Identify LED is on.
   - If you don't know the physical location based on the Chassis name, MTMS, or frame, you will have to walk the floor and look for Identify LEDs that are on and then verify that the switch at which you are looking matches the MTMS for the FRU.

How to find an adapter FRU using another device with a valid location code

How to find an adapter FRU using another device with a valid location code

Sometimes there is a problem with the location code of a FRU. In such cases, you may be able to use the location code of a FRU to which the desired FRU is attached. This works under the assumption that none of the other FRUs in the same chassis with the FRU with the invalid location will have valid location codes.

It is possible that the location code for the desired FRU is now available in the IBM Network Manager.

Note: For this procedure to work, the network should not have been re-cabled since the time of the latest occurrence of the serviceable event. It is particular important that the failing link was not re-cabled.

1. Find and record information:
   a. Record the location code for the known device.
   b. If the known device has port information in the location code, proceed to the next step. Otherwise, you will need to go to "How to find an adapter FRU using a reference code extension" on page 176.
   c. Proceed only if you have a known device with port information.
      1) Open the IBM Network Manager to the View Switch Topology window.
      2) Using its location code, find the known device in the window.
      3) Record the Neighbor Location-Code, Neighbor Name, and Neighbor Frame:Cage.
      4) If there is no neighbor information, it is possible that the neighbor is a PCI-X adapter that is not currently supported by the IBM Network Manager. If that is the case, you will either have to refer to cable labels, or cable planning documentation to determine the location of the PCI-X adapter. **This procedure ends here.**
      5) If there is neighbor information and the neighbor is a switch, **this procedure ends here.**
      6) If there is neighbor information and the neighbor is an adapter, proceed to the next step with that recorded information.
   d. You have determined that the known device is an adapter port, so perform the following procedure to find the server that contains the adapter:
      1) Record the location code of the known device
      2) Open the IBM Network Manager to the View End-Point window.
      3) Using its location code, find the known device in the window.
      4) Record the Neighbor Location-Code, Neighbor Name, and Neighbor Frame:Cage.
2. Turn on Identify LEDs:
   a. Go to the controlling HMC for the server.
   b. g. Using Service Focal Point-Service Utilities, turn on the identify LED for the server.
3. Go to physical location and verify FRU:
   - If you know the physical location of the desired switch by the name, the MTMS, or the frame, go to that switch, and verify that the Identify LED is on.
   - If you don't know the physical location based on the Chassis name, MTMS, or frame, you will have to walk the floor and look for Identify LEDs that are on and then verify that the switch at which you are looking matches the MTMS for the FRU.

How to find a switch FRU using a reference code extension

How to find a switch FRU using a reference code extension.

In some cases, there are no valid location codes in a FRU list. In such cases, you will need to work with a logical location id which is recorded in the reference code extension in SFP. Everything after the U[*machine type*].[*model*].[*serial number*] portion of the location code should be valid, so the key to this procedure is to determine frame and chassis number of the device, because that will get you to the switch chassis. The remainder of the location code should then match a device within that chassis.

Note: For this procedure to work, the network should not have been re-cabled since the time of the latest occurrence of the serviceable event. It is particular important that the failing link was not re-cabled.

1. Find and record information about the FRU:
   a. Record the location code for the FRU. Everything except for the U[*MachineType*].[*model*].[*serial number*] portion of the location code should be valid.
   b. Check the other FRUs. Does another FRU have the same location code? If one does, then it will be impossible to determine if this FRU is the one that matches the reference code extension. This should not be an issue, because in such instances, there is equal probability that either side is the failing device.
   c. Record the reference code extension from the serviceable event.
   d. Record the frame, card number and port number from the reference code extension:

Network	**Frame**	**Chassis**	**Chassis Type**	**Device Type**	**Card**	**Chip**	**Port**
4 Hex character	3 Hex character	2 Hex character	Hex character	Hex character	2 Hex characters	Hex character	2 Hex characters

   If there are fewer than 16 characters, then leading zeroes have been dropped, for example: 0001002052301100

   network = 1; frame=2; cage=5; cage type=5; device type=3; card=01; chip=1; port=00
2. Determine if reference code extension applies to this FRU or the other side of the link
   a. If above you determined that you can't tell which side is which, then we will assume that the desired FRU is described in the reference code extension and proceed to step 3. Otherwise proceed to the next step.
   b. Match the card and port numbers to their corresponding numbers in the location code: U[*machinetype*].[*model*].[*serial number*]-P[*planar*]-C[*card number*]-T[*port number*]
   c. If the card number and port number match, then the desired FRU is described in the reference code extension. Record this and proceed to step 3.
   d. If you have arrived here, the reference code extension describes the device attached to the desired FRU. Record this and proceed to step 3.
3. Determine the frame number of the desired FRU
   a. If the reference code extension describes the desired FRU, record the frame number in the reference code extension. (Do not add one to it as you may have done with the card and port numbers.) Proceed to Turn on the identify LED.
   b. If the reference code extension describes the device attached to the desired FRU, perform the following to determine the frame number of the desired FRU. Otherwise skip to the next step. (You will note that because only switches report link errors through the IBM Network Manager, the reference code extension will always relate to a switch.)
      1) Open the IBM Network Manager to the View Switch Topology window.
      2) Find the device described by the location code and frame number you recorded above. The frame number will act in place of the unit location field: U[*MachineType*].[*model*].[*serial number*].
      3) Record the Neighbor Location-Code, Neighbor Name and Neighbor Frame:Cage. These are the location code, chassis name and frame of the desired FRU.
4. Turn on the Identify LED
   a. Go to the View Switch Topology window and find the desired FRU using the location code, chassis name and frame recorded in previous steps.
   b. On the menu in the View Switch Topology window, choose **Selected-Identify-On**.
5. Go to physical location and verify FRU.
   - If you know the physical location of the switch by the Chassis name, the MTMS, or the frame, go to that switch, and verify that the Identify LED is on.

- If you don't know the physical location based on the Chassis name, MTMS, or frame, you will have to walk the floor and look for Identify LEDs that are on and then verify that the switch at which you are looking matches the MTMS for the FRU.

How to find an adapter FRU using a reference code extension

How to find an adapter FRU using a reference code extension.

In some cases, there are no valid location codes in a FRU list. In such cases, you will need to work with a logical location id which is recorded in the reference code extension in SFP. Because the IBM Network Manager only presents serviceable events that come from switches, the Reference Code Extension will always be for the switch attached to the desired adapter FRU. So, the basic approach for this procedure is to first find the switch port that reported the serviceable event, and then find the adapter on the other side of the cable.

Note: For this procedure to work, the network should not have been re-cabled since the time of the latest occurrence of the serviceable event. It is particular important that the failing link was not re-cabled.

1. Record Service Focal Point Information:
   a. Record the location code for the attached switch FRU which is also in the FRU list. Everything except for the U[*MachineType*].[*model*].[*serial number*] portion of the location code should be valid.
   b. Record the reference code extension from the serviceable event.
   c. Record the frame, card number and port number from the reference code extension:

Network	Frame	Chassis	Chassis Type	Device Type	Card	Chip	Port
4 Hex character	3 Hex character	2 Hex character	Hex character	Hex character	2 Hex characters	Hex character	2 Hex characters

   If there are fewer than 16 characters, then leading zeroes have been dropped, for example: 0001002052301100

   network = 1; frame=2; cage=5; cage type=5; device type=3; card=01; chip=1; port=00
2. Find and record information about the adapter:
   a. Go to the IBM Network Manager View Switch Topology window.
   b. Using the frame and location code for the attached switch recorded above, find the attached switch.
   c. Record the Neighbor Location-Code, Neighbor Name and Neighbor Frame:Cage. These are the location code, name and frame and cage of the desired FRU. It is possible that the location code found in the View Switch Topology window is now valid.
3. Turn on the Identify LED
   a. Go to the controlling HMC for the managed system. The managed system should be identified by one of the following methods based on information gathered about the adapter:
      1) If you found a valid location code, the managed server is identified in the unit location field U[*FeatureCode*].001.[*SerialNumber*]-.
      2) If this is a high-end server in a 24-inch frame, the frame and cage numbers will help identify the location of the server.
      3) The name of the chassis can identify the server.
   b. Using Service Focal Point-Service Utilities turn on the Identify LED for the managed system.
4. Go to physical location and verify FRU.
   - If you know the physical location of the switch by the Chassis name, the MTMS, or the frame, go to that switch, and verify that the Identify LED is on.

- If you don't know the physical location based on the Chassis name, MTMS, or frame, you will have to walk the floor and look for Identify LEDs that are on and then verify that the switch at which you are looking matches the MTMS for the FRU.

How to determine a switch FRU using the error description

Put your short description here; used for first paragraph and abstract.

For card level and chassis level FRUs it may be possible to use the error description in Service Focal Point to determine the location of the FRU. The frame and chassis/cage information is the most important; however, the slot information can also be useful. This procedure should only be used if only card and chassis level FRUs are being called out.

1. Determine if this is the correct procedure:
   a. If you have another switch FRU listed with a valid location code, you should use the procedure "How to find a switch FRU using another device with a valid location code" on page 223, and stop using this procedure.
   b. If you do not have any port level or cable FRUs, which contain –T# in their location code, then the error description has information pertinent to the desired FRU, and you should continue on to Record Service Focal Point Information.
   c. If there is port level or cable FRU in the error description, you will be better served by using the procedure, "How to find an adapter FRU using a reference code extension" on page 176.
2. Record Service Focal Point Information
   a. Record the frame, cage, and slot information in the error description in SFP. Remember that the cage and slot are not available in all error descriptions.
   b. Record the location code for the desired FRU. The frame and chassis information recorded above will help to find the chassis when the unit location (U[*MachineType*].[*model*].[*SerialNumber*]) portion of the location code is not valid.
3. Turn on the Identify LED
   a. Go to the View Switch Topology window and find the desired FRU using the location code and frame recorded in previous steps. Record the chassis name.
   b. On the menu in the View Switch Topology window, choose **Selected-Identify-On**
4. Go to physical location and verify FRU:
   - If you know the physical location of the switch by the Chassis name, the MTMS, or the frame, go to that switch, and verify that the Identify LED is on.
   - If you don't know the physical location based on the Chassis name, MTMS, or frame, you will have to walk the floor and look for Identify LEDs that are on and then verify that the switch at which you are looking matches the MTMS for the FRU.

How to determine an adapter FRU using the error description

How to Determine an adapter FRU using the error description

For card level and chassis level FRUs it may be possible to use the error description in Service Focal Point to determine the location of the FRU. The frame and chassis information is the most important; however, the slot information can also be useful.

Because you should never have a situation where an adapter is the only FRU in the FRU list for a serviceable event that is sourced by the IBM Network Manager, you should use one of the two procedures listed below:

1. If you have a switch FRU listed with a valid location code, you should use the procedure, "How to find an adapter FRU using another device with a valid location code" on page 223, and stop using this procedure.

2. If there is a port level or cable FRU which is indicated with a location code of the format: U[*value*].[*value*].[*value*]-P*x*-C*y*-T*z*, you will be better served by using the procedure, "How to find an adapter FRU using a reference code extension" on page 176, and stop using this procedure.

InfiniBand parts information

See the following sections for parts information.

How to determine an unknown part number for an InfiniBand switch network component

How to determine an unknown part number for an InfiniBand switch network component

Note: IBM Service responsibilities include the cable and HCAs. For SFS7000P and SFS7008P InfiniBand switches, IBM Service responsibilities do not include the InfiniBand switches, nor any FRU, software, or firmware contained within them. If the switch is a 7048-120 or 7048-270 machine type, IBM Service has responsibility for it.

To determine the part number of a switch component you may do one of the following:

1. Look up the part number in the parts catalog.
   a. If this is a Topspin switch component, consider the following before looking up the part number.
      1) A 7048-120 and SFS7000P only have three FRUs:
         a) Switch chassis, which covers all but the FRUs listed below
         b) Combined fan and power card
         c) Battery
      2) A 7048-270 and SFS7008P have many more potential FRUs than the 7048-120.
      3) The Topspin switch catalog can be found here: “Parts catalog”
   b. If this is an adapter, look in the parts catalog for the type of server in which it is installed.
2. Remove the part and look at the part number label.

Identifying InfiniBand cables and cabling expansion units

For information on identifying InfiniBand cables and cabling expansion units, see the following topics in the in the IBM Systems Hardware Information Center

Expansion unit

Remote I/O, High speed link or InfiniBand adapter

Removing and replacing InfiniBand parts

Refer to the Topspin documentation for InfiniBand switch service procedures.

Parts catalog

Table 24. Adapter and cable FRU parts

IBM FRU Part Number	MTM	Specify code	Description
03N6028			GX Dual-port 4x IB HCA (1809)
39J2706			GX Dual-port 4x IB HCA bracket
39J2069			GX Dual-port 4x IB HCA (1810)
41U0383			GX Dual-port 4x IB HCA (1811)

© Copyright IBM Corp. 2004, 2007

Table 24. Adapter and cable FRU parts (continued)

IBM FRU Part Number	MTM	Specify code	Description
16R0802			Filler for GX Dual-port 4x IB HCA
39J2444			GX dual-port 4x IB HCA (1812)
39J2623			Cable support for GX dual-port 4x IB HCA (1812)
97P6113			3 meter 4x IB Cable (1835)
21P8318			8 meter 4x IB Cable (1836)
21P8319			8 meter 12x to three 4x IB Cable (1838)
39J4319			1.5 meter 4x IB Cable (1839)
60H1866			GX Dual-port 12x IB HCA (7820)
12R7678			Filler for GX Dual-port 12x IB HCA (7820)

Table 25. Switch FRU parts

IBM FRU Part Number	MTM	Specify code	Description
41V0246	7048-120		pSeries VPD 120 switch chassis FRU
41V0245	7048-120		Topspin 120, power supply/cooling unit
41V0248	7048-270		Topspin 270 switch chassis FRU
26K7538	7048-270		Topspin 270 management interface module
26K7539	7048-270		Topspin 270 fan tray
26K7540	7048-270		Topspin 270 power supply
26K7541	7048-270		Topspin 270 rail kit
26K7542	7048-270		Topspin 270 serial cable kit
26K7543	7048-270		Topspin 270 chassis ID card
26K7544	7048-270		Topspin 270 blanking panels (2x LIM SFM)
26K7545	7048-270		Topspin 270 12-port 4x line interface module
41V0247	7048-270		Topspin 270 4-port 12x line interface module
26K7546	7048-270		Topspin 270 switch fabric module
26K7629	7048-270		Topspin 270 bezel
39M5080	7048-120		Topspin 120 power cord
41V0794	7048-120, 7048-270		IBM 7014 rack switch mounting fasteners • 12 nutclips • 12 mounting screws

Appendix A. Notices

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

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

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

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: THIS INFORMATION IS PROVIDED "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. The manufacturer 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 Web sites not owned by the manufacturer are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this product and use of those Web sites is at your own risk.

The manufacturer 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 products not produced by this manufacturer was obtained from the suppliers of those products, their published announcements or other publicly available sources. This manufacturer has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to products not produced by this manufacturer. Questions on the capabilities of products not produced by this manufacturer should be addressed to the suppliers of those products.

All statements regarding the manufacturer's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

© Copyright IBM Corp. 2004, 2007

The manufacturer's prices shown are the manufacturer's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

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.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

The drawings and specifications contained herein shall not be reproduced in whole or in part without the written permission of the manufacturer.

The manufacturer has prepared this information for use with the specific machines indicated. The manufacturer makes no representations that it is suitable for any other purpose.

The manufacturer's computer systems contain mechanisms designed to reduce the possibility of undetected data corruption or loss. This risk, however, cannot be eliminated. Users who experience unplanned outages, system failures, power fluctuations or outages, or component failures must verify the accuracy of operations performed and data saved or transmitted by the system at or near the time of the outage or failure. In addition, users must establish procedures to ensure that there is independent data verification before relying on such data in sensitive or critical operations. Users should periodically check the manufacturer's support websites for updated information and fixes applicable to the system and related software.

Trademarks

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

AIX
AIX 5L
Domino
POWER
POWER5
400
AS/400
DB2
DB2 Universal Database
e(logo)server
Electronic Service Agent
Enterprise Storage Server
eServer
HACMP
i5/OS
IBM
IBM Workplace
iSeries
Lotus
Midrange Express
NetServer
NetView
OS/400

OpenPower
pSeries
PTX
Rational
Redbooks
RS/6000
ServicePac
Tivoli
TotalStorage
Virtualization Engine
VTAM
WebSphere
xSeries

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

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

Regulatory notices

Class A Notices - Federal Communications Commission (FCC) statement

Note: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Responsible Party:

International Business Machines Corporation
New Orchard Road
Armonk, NY 10504

Telephone: 1-919-543-2193

Industry Canada Compliance Statement

This Class A digital apparatus meets the requirements of the Canadian Interference-Causing Equipment Regulations.

Avis de conformité à la réglementation d'Industrie Canada

Cet appareil numérique de la classe A respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

European Community Compliance Statement

This product is in conformity with the protection requirements of EU Council Directive 89/336/EEC on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

This product has been tested and found to comply with the limits for Class A Information Technology Equipment according to European Standard EN 55022. The limits for Class A equipment were derived for commercial and industrial environments to provide reasonable protection against interference with licensed communication equipment.

Warning: This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

VCCI Statement - Japan

この装置は、情報処理装置等電波障害自主規制協議会（ＶＣＣＩ）の基準に基づくクラスＡ情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

The following is a summary of the VCCI Japanese statement in the box above.

This is a Class A product based on the standard of the Voluntary Control Council for Interference by Information Technology Equipment (VCCI). If this equipment is used in a domestic environment, radio disturbance may arise. When such trouble occurs, the user may be required to take corrective actions.

Electromagnetic Interference (EMI) Statement - People's Republic of China

声　明

此为Ａ级产品，在生活环境中，
该产品可能会造成无线电干扰。
在这种情况下，可能需要用户对其
干扰采取切实可行的措施。

Declaration: This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may need to perform practical action.

Electromagnetic Interference (EMI) Statement - Taiwan

警告使用者：
這是甲類的資訊產品，在
居住的環境中使用時，可
能會造成射頻干擾，在這
種情況下，使用者會被要
求採取某些適當的對策。

The following is a summary of the EMI Taiwan statement above.

Warning: This is a Class A product. In a domestic environment this product may cause radio interference in which case the user will be required to take adequate measures.

IBM Taiwan Contact Information

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

Electromagnetic Interference (EMI) Statement - Korea

이 기기는 업무용으로 전자파적합등록을 한 기기이오니 판매자 또는 사용자는 이점을 주의하시기 바라며, 만약 잘못 판매 또는 구입하였을 때에는 가정용으로 교환하시기 바랍니다.

Radio Protection for Germany

Dieses Gerät ist berechtigt in Übereinstimmung mit Dem deutschen EMVG vom 9.Nov.92 das EG–Konformitätszeichen zu führen.

Der Aussteller der Konformitätserklärung ist die IBM Germany.

Dieses Gerät erfüllt die Bedingungen der EN 55022 Klasse A. Für diese von Geräten gilt folgende Bestimmung nach dem EMVG:

Geräte dürfen an Orten, für die sie nicht ausreichend entstört sind, nur mit besonderer Genehmigung des Bundesministers für Post und Telekommunikation oder des Bundesamtes für Post und Telekommunikation betrieben werden. Die Genehmigung wird erteilt, wenn keine elektromagnetischen Störungen zu erwarten sind.

(Auszug aus dem EMVG vom 9.Nov.92, Para.3, Abs.4)

Hinweis

Dieses Genehmigungsverfahren ist von der Deutschen Bundespost noch nicht veröffentlicht worden.

Class B Notices - Federal Communications Commission (FCC) statement

Note: This equipment has been tested and found to comply with the limits for a class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.

- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult an IBM authorized dealer or service representative for help.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. Proper cables and connectors are available from IBM authorized dealers. IBM is not responsible for any radio or television interference caused by using other than recommended cables or connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interferences, and (2) this device must accept any interferences received, including interference that may cause undesired operation.

Responsible Party:

International Business Machines Corporation
New Orchard Road
Armonk, NY 10504

Telephone: 1-919-543-2193

Industry Canada Compliance Statement

This Class B digital apparatus meets the requirements of the Canadian Interference-Causing Equipment Regulations.

Avis de conformité à la réglementation d'Industrie Canada

Cet appareil numérique de la classe B respecte toutes les exigences du Réglement sur le matériel brouilleur du Canada.

European Community Compliance Statement

This product is in conformity with the protection requirements of EC Council Directive 89/336/EEC on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

This product has been tested and found to comply with the limits for Class B Information Technology Equipment according to CISPR 22 / European Standard EN 55022. The limits for Class B equipment were derived for typical residential environments to provide reasonable protection against interference with licensed communication devices.

Properly shielded and grounded cables and connectors must be used in order to reduce the potential for causing interference to radio and TV communications and to other electrical or electronic equipment. Such cables and connectors are available from IBM authorized dealers. IBM cannot accept responsibility for an interference caused by using other than recommended cables and connectors.

European Community contact:
IBM Technical Regulations
Pascalstr. 100, Stuttgart, Germany 70569
Tele: 0049 (0)711 785 1176
Fax: 0049 (0)711 785 1283
E-mail: tjahn@de.ibm.com

VCCI Statement - Japan

この装置は、情報処理装置等電波障害自主規制協議会（ＶＣＣＩ）の基準に基づくクラスＢ情報技術装置です。この装置は、家庭環境で使用することを目的としていますが、この装置がラジオやテレビジョン受信機に近接して使用されると、受信障害を引き起こすことがあります。
取扱説明書に従って正しい取り扱いをして下さい。

The following is a summary of the VCCI Japanese statement in the box above.

This is a Class B product based on the standard of the Voluntary Control Council for Interference by Information Technology Equipment (VCCI). If this equipment is used in a domestic environment, radio disturbance may arise. When such trouble occurs, the user may be required to take corrective actions.

IBM Taiwan Product Service Contact Information

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

Electromagnetic Interference (EMI) Statement - Korea

이 기기는 가정용으로 전자파적합등록을 한 기기로서 주거 지역에서는 물론 모든 지역에서 사용할 수 있습니다.

Radio Protection for Germany

Dieses Gerät ist berechtigt in Übereinstimmung mit Dem deutschen EMVG vom 9.Nov.92 das EG–Konformitätszeichen zu führen.

Der Aussteller der Konformitätserklärung ist die IBM Germany.

Dieses Gerät erfüllt die Bedingungen der EN 55022 Klasse B. Für diese von Geräten gilt folgende Bestimmung nach dem EMVG:

Geräte dürfen an Orten, für die sie nicht ausreichend entstört sind, nur mit besonderer Genehmigung des Bundesministers für Post und Telekommunikation oder des Bundesamtes für Post und Telekommunikation betrieben werden. Die Genehmigung wird erteilt, wenn keine elektromagnetischen Störungen zu erwarten sind.

(Auszug aus dem EMVG vom 9.Nov.92, Para.3, Abs.4)

Hinweis

Dieses Genehmigungsverfahren ist von der Deutschen Bundespost noch nicht veröffentlicht worden.

Terms and conditions

Permissions for the use of these publications is granted subject to the following terms and conditions.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of the manufacturer.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of the manufacturer.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any data, software or other intellectual property contained therein.

The manufacturer reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by the manufacturer, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

THE MANUFACTURER MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THESE PUBLICATIONS ARE PROVIDED ″AS-IS″ AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Product recycling and disposal

This unit must be recycled or discarded according to applicable local and national regulations. IBM encourages owners of information technology (IT) equipment to responsibly recycle their equipment when it is no longer needed. IBM offers a variety of product return programs and services in several countries to assist equipment owners in recycling their IT products. Information on IBM product recycling offerings can be found on IBM's Internet site at http://www.ibm.com/ibm/environment/products/prp.shtml.

Note: This mark applies only to countries within the European Union (EU) and Norway.

Appliances are labeled in accordance with European Directive 2002/96/EC concerning waste electrical and electronic equipment (WEEE). The Directive determines the framework for the return and recycling of used appliances as applicable throughout the European Union. This label is applied to various products to indicate that the product is not to be thrown away, but rather reclaimed upon end of life per this Directive.

In accordance with the European WEEE Directive, electrical and electronic equipment (EEE) is to be collected separately and to be reused, recycled, or recovered at end of life. Users of EEE with the WEEE

marking per Annex IV of the WEEE Directive, as shown above, must not dispose of end of life EEE as unsorted municipal waste, but use the collection framework available to customers for the return, recycling, and recovery of WEEE. Customer participation is important to minimize any potential effects of EEE on the environment and human health due to the potential presence of hazardous substances in EEE. For proper collection and treatment, contact your local IBM representative.

Battery return program

This product may contain sealed lead acid, nickel cadmium, nickel metal hydride, lithium, or lithium ion battery. Consult your user manual or service manual for specific battery information. The battery must be recycled or disposed of properly. Recycling facilities may not be available in your area. For information on disposal of batteries outside the United States, go to http://www.ibm.com/ibm/environment/products/batteryrecycle.shtml or contact your local waste disposal facility.

In the United States, IBM has established a return process for reuse, recycling, or proper disposal of used IBM sealed lead acid, nickel cadmium, nickel metal hydride, and other battery packs from IBM Equipment. For information on proper disposal of these batteries, contact IBM at 1-800-426-4333. Please have the IBM part number listed on the battery available prior to your call.

For Taiwan: Please recycle batteries.

IBM Cryptographic Coprocessor Card Return Program

The following information applies only for systems originally sold prior to July 1, 2006:

This machine may contain an optional feature, the cryptographic coprocessor card, which includes a polyurethane material that contains mercury. Please follow local ordinances or regulations for disposal of this card. IBM has established a return program for certain IBM Cryptographic Coprocessor Cards. More information can be found at http://www.ibm.com/ibm/environment/products/prp.shtml.

Appendix B. Planning and Installation Worksheets

The worksheets in this section can be copied and used when planning and installing a cluster.

Installation coordinating worksheets

Worksheets to help plan and coordinate installation tasks for an InfiniBand cluster and the IBM Network Manager.

Copy the following worksheets to help coordinate installation tasks. Each organization should use a separate installation sheet and it should be completed using the flow illustrated in Figure 3 on page 15.

Table 26. Installation coordinating worksheet

Organization:				
Task	**Task Description**	**Prerequisite tasks**	**Scheduled date**	**Completed date**

© Copyright IBM Corp. 2004, 2007

Planning checklist

The planning checklist helps you track your progress through the planning process.

Table 27. Planning checklist

Step	Target	Complete
Start planning checklist		
Gather documentation and review planning information for individual devices.		
Ensure that you have planned for: • Servers • I/O devices • InfiniBand network devices • Frames or racks • Service network, including: – HMCs – Ethernet devices – CSM Management Server (for multiple HMC environments) – AIX NIM server (for servers with no removable media) – Linux distribution server (for servers with no removable media) • System management applications (HMC and CSM and IBM Network Manager) • Physical dimension and weight characteristics • Electrical characteristics • Cooling characteristics		
Ensure that you have the "Required levels of support, firmware, and devices" on page 7 for your network hardware		
Review cabling and topology documentation for InfiniBand networks		
Review "Installation flow of clusters using the IBM Network Manager" on page 11		
Review "Planning for a high-performance computing message-passing interface configuration" on page 18		
Review "Planning octopus cables in static 12x cabling" on page 19		
Review "Planning Aids" on page 22		
Complete planning worksheets		
Complete planning process		
Review readme files and online information related to the software and firmware to ensure that you have up-to-date information and the latest supported levels		

Planning worksheets

The planning worksheets are used when you are planning your cluster.

Tip: It is best to keep the sheets somewhere that is accessible to the system administrators and service representatives not only during the installation process, but also for future reference during maintenance, upgrade, or repair actions.

Cluster summary worksheet

Record on the items which

Cluster summary worksheet
Cluster name:
Application:
Number and types of servers:
Number of servers and HCAs per server: **Note:** If there are servers with various numbers of HCAs, list the number of servers with each configuration; for example, 12 servers with one 2-port HCA; 4 servers with two 2-port HCAs.
Number of 7048-120 or SFS7000P switches:
Number of 7048-270 or SFS7008P switches:
Number of subnets:
List of GID-prefixes and subnet masters (assign a number to a subnet for easy reference):
Switch partitions:
Number of frames:
Number of HMCs:
CSM and Cluster Ready Hardware Server used?
Number of Service Ethernet networks:
Service network domains:
Service network DHCP server locations:
Service network switches with static IP:
Service network HMCs with static IP:
Service network DHCP range(s):
AIX NIM server info:
Linux distribution server info:
Power requirements:
Maximum cooling required:
Number of cooling zones:
Maximum weight per square foot:

Frame and rack planning worksheet

The frame and rack planning worksheet is used for planning how to populate frames or racks. You will likely have to understand the numbers of each device type (servers, switches, and bulk power adapters). For the slots, indicate the range of slots or drawers that a device will populate. A standard method for naming slots can either be found in the documentation for the frames or servers, or use EIA heights (1.75 inches) as a standard.

Frame and rack planning worksheet		
Frame or rack number(s): _____ Frame MTM or feature: _____ Frame size: _____ (19-inch or 24-inch) Number of slots: _____		
Slots	**Device type (server, switch, or BPA,) Indicate Machine type and model**	**Device name**

Server planning worksheet

You can use this worksheet as a template for multiple servers with similar configurations. For such cases, you will want to give the range of names of these servers and where they are located. Also, you can use the configuration note to remind you of other specific characteristics of the server. It is important to note the type of HCAs to be used.

Server planning worksheet				
Name(s): _____ Type(s): _____ Frame(s)/slot(s): _____ Number and type of HCAs_____ Num LPARs/LHCAs: _____ IP addressing for InfiniBand: _____ IP addressing of service subsystem: _____ LPAR IP addressing: _____ MPI addressing: _____ Configuration notes:				
HCA information				
HCA	Capability (Sharing)	HCA port	Switch connection	GID prefix
LPAR information				
LPAR/LHCA (give name)	OS Type	GUID index	Shared HCA (capability)	Switch partition

Switch planning worksheet

There is a worksheet in this section for each type of switch.

Use the following worksheet for planning 7048-270 and SFS7008P switches.

7048-270 and SFS7008P switch worksheet	
Switch MTM: _____ (7048-270 or SFS7008P) Switch name: _____ Frame and slot: _____ IP address: _____ (if single HMC in cluster, indicate DHCP) GID-prefix: _____ LMC: _____ (0=default; 2=if used in HPC cluster) Switch MTMS: _____ (Complete during installation)	
Ports	**Connection**
1 (16)	
2	
3	
4 (16)	
5	
6	
7 (16)	
8	
9	
10 (16)	
11	
12	
13	
14	
15 (16)	
16	
17	
18 (16)	
19	
20	
21 (16)	
22	
23	
24 (16)	

Use the following worksheet for planning 7048-270 and SFS7008P switches.

7048-270 and SFS7008P switch worksheet			
Switch MTM: _____ (7048-270 or SFS7008P) Switch name: _____ Frame and slot: _____ IP address: _____ (if single HMC in cluster, indicate DHCP) GID-prefix: _____ LMC: _____ (0=default; 2=if used in HPC cluster) Switch MTMS: _____ (Complete during installation)			
LIM 1		**LIM 2**	
Ports	**Connection**	**Ports**	**Connection**
1		1 (16)	
2		2	
3 (16)		3	
4		4 (16)	
5		5	
6 (16)		6	
7 (16)		7	
8		8	
9		9 (16)	
10 (16)		10	
11		11	
12		12 (16)	
LIM 3		**LIM 4**	
Ports	**Connection**	**Ports**	**Connection**
1 (16)		1	
2		2	
3		3 (16)	
4 (16)		4	
5		5	
6		6 (16)	
7		7 (16)	
8		8	
9 (16)		9	
10		10 (16)	
11		11	
12 (16)		12	
LIM 5		**LIM 6**	
Ports	**Connection**	**Ports**	**Connection**
1		1 (16)	
2		2	
3 (16)		3	
4		4 (16)	

7048-270 and SFS7008P switch worksheet			
5		5	
6 (16)		6	
7 (16)		7	
8		8	
9		9 (16)	
10 (16)		10	
11		11	
12		12 (16)	
LIM 7		**LIM 8**	
Ports	**Connection**	**Ports**	**Connection**
1 (16)		1	
2		2	
3		3 (16)	
4 (16)		4	
5		5	
6		6 (16)	
7		7 (16)	
8		8	
9 (16)		9	
10		10 (16)	
11		11	
12 (16)		12	

Printed in USA