IBM

eserver

IBM eserver Cluster 1350

Installation and Service Guide

IBM eserver Cluster 1350

Installation and Service Guide

Note: Before using this information and the product it supports, read the general information in “Safety” on page ix and Appendix F, “Notices,” on page 103.

Sixth edition, June 2004

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

Contents

© Copyright IBM Corp. 2004

Safety

For general information concerning safety, refer to *Electrical Safety for IBM Customer Engineers*, S229-8124. For a copy of the publication, contact your IBM$^{®}$ account representative or the IBM branch office serving your locality.

Enterprise rack safety information: Read the safety notices in the manual provided with the enterprise rack before beginning work. Keep the Enterprise Rack manual near the rack for fast reference.

The procedures described in this document must be performed by qualified service personnel. Safety warnings are contained within these procedures. If you cannot read the language of this document, do not perform any procedures until you receive a translated copy. IBM does not accept responsibility or liability for failure to follow these procedures correctly.

Safety Information

Before installing this product, read the Safety Information.

قبل تركيب هذا المنتج، يجب قراءة الملاحظات الآمنية

Antes de instalar este produto, leia as Informações de Segurança.

在安装本产品之前，请仔细阅读 Safety Information
（安全信息）。

安裝本產品之前，請先閱讀「安全資訊」。

Prije instalacije ovog produkta obavezno pročitajte Sigurnosne Upute.

Před instalací tohoto produktu si přečtěte příručku bezpečnostních instrukcí.

Læs sikkerhedsforskrifterne, før du installerer dette produkt.

Lees voordat u dit product installeert eerst de veiligheidsvoorschriften.

Ennen kuin asennat tämän tuotteen, lue turvaohjeet kohdasta Safety Information.

Avant d'installer ce produit, lisez les consignes de sécurité.

Vor der Installation dieses Produkts die Sicherheitshinweise lesen.

Πριν εγκαταστήσετε το προϊόν αυτό, διαβάστε τις πληροφορίες ασφάλειας (safety information).

לפני שתתקינו מוצר זה, קראו את הוראות הבטיחות.

A termék telepítése előtt olvassa el a Biztonsági előírásokat!

Prima di installare questo prodotto, leggere le Informazioni sulla Sicurezza.

製品の設置の前に、安全情報をお読みください。

본 제품을 설치하기 전에 안전 정보를 읽으십시오.

Пред да се инсталира овој продукт, прочитајте информацијата за безбедност.

© Copyright IBM Corp. 2004

Les sikkerhetsinformasjonen (Safety Information) før du installerer dette produktet.

Przed zainstalowaniem tego produktu, należy zapoznać się
z książką "Informacje dotyczące bezpieczeństwa" (Safety Information).

Antes de instalar este produto, leia as Informações sobre Segurança.

Перед установкой продукта прочтите инструкции по
технике безопасности.

Pred inštaláciou tohto zariadenia si pečítaje Bezpečnostné predpisy.

Pred namestitvijo tega proizvoda preberite Varnostne informacije.

Antes de instalar este producto, lea la información de seguridad.

Läs säkerhetsinformationen innan du installerar den här produkten.

Important:

> All caution and danger statements in this documentation begin with a number. This number is used to cross reference an English caution or danger statement with translated versions of the caution or danger statement in the *IBM NetBAY Rack Safety Information* book.
>
> For example, if a caution statement begins with a number 1, translations for that caution statement appear in the *IBM NetBAY Rack Safety Information* book under statement 1.
>
> Be sure to read all caution and danger statements in this documentation before performing the instructions. Read any additional safety information that comes with your server or optional device before you install the device.

Statement 2:

⚠

DANGER

- **Always lower the leveling pads on the rack cabinet.**
- **Always install stabilizer brackets on the rack cabinet.**
- **Always install servers and optional devices starting from the bottom of the rack cabinet.**
- **Always install the heaviest devices in the bottom of the rack cabinet.**

Statement 3:

DANGER

- **Do not extend more than one sliding device at a time.**
- **The maximum allowable weight for devices on slide rails is 80 kg (176 lb). Do not install sliding devices that exceed this weight.**

Class 1 Laser Product
Laser Klasse 1
Laser Klass 1
Luokan 1 Laserlaite
Appareil À Laser de Classe 1

Statement 4:

DANGER

Electrical current from power, telephone, and communication cables is hazardous.

To avoid a shock hazard:

- **Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.**
- **Connect all power cords to a properly wired and grounded electrical outlet.**
- **Connect to properly wired outlets any equipment that will be attached to this product.**
- **When possible, use one hand only to connect or disconnect signal cables.**
- **Never turn on any equipment when there is evidence of fire, water, or structural damage.**
- **Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.**
- **Connect and disconnect cables as described in the following table when installing, moving, or opening covers on this product or attached devices.**

To Connect:	To Disconnect:
1. Turn everything OFF. 2. First, attach all cables to devices. 3. Attach signal cables to connectors. 4. Attach power cords to outlet. 5. Turn device ON.	1. Turn everything OFF. 2. First, remove power cords from outlet. 3. Remove signal cables from connectors. 4. Remove all cables from devices.

Statement 5:

CAUTION:
The power control button on the device and the power switch on the power supply do not turn off the electrical current supplied to the device. The device also might have more than one power cord. To remove all electrical current from the device, ensure that all power cords are disconnected from the power source.

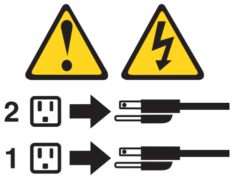

Statement 6:

CAUTION:
If you install a strain-relief bracket option over the end of the power cord that is connected to the device, you must connect the other end of the power cord to an easily accessible power source.

Statement 7:

CAUTION:
If the device has doors, be sure to remove or secure the doors before moving or lifting the device to avoid personal injury. The doors will not support the weight of the device.

Statement 8:

DANGER

- **Plug power cords from devices in the rack cabinet into electrical outlets that are located near the rack cabinet and are easily accessible.**
- **Each rack cabinet might have more than one power cord. Be sure to disconnect all power cords in the rack cabinet before servicing any device in the rack cabinet.**
- **Install an emergency-power-off switch if more than one power device (power distribution unit or uninterruptible power supply) is installed in the same rack cabinet.**
- **Connect all devices installed in a rack cabinet to power devices installed in the same rack cabinet. Do not plug a power cord from a device installed in one rack cabinet into a power device installed in a different rack cabinet.**

Statement 10:

CAUTION:
Removing components from the upper positions in the Enterprise Rack cabinet improves rack stability during relocation. Follow these general guidelines whenever you relocate a populated rack cabinet within a room or building:

- **Reduce the weight of the rack cabinet by removing equipment starting at the top of the rack cabinet. When possible, restore the rack cabinet to the configuration of the rack cabinet as you received it. If this configuration is not known, you must do the following:**
  - **Remove all devices in the 32U position and above.**
  - **Ensure that the heaviest devices are installed in the bottom of the rack cabinet.**
  - **Ensure that there are no empty U-levels between devices installed in the rack cabinet below the 32U level.**
- **If the rack cabinet you are relocating is part of a suite of rack cabinets, detach the rack cabinet from the suite.**
- **Inspect the route that you plan to take to eliminate potential hazards.**
- **Make sure that the route that you choose can support the weight of the loaded rack cabinet. Refer to the documentation that comes with your rack cabinet for the weight of a loaded rack cabinet.**
- **Make sure that all door openings are at least 760 x 2030 MM. (30 x 80 in.)**
- **Ensure that all devices, shelves, drawers, doors, and cables are secure.**
- **Ensure that the four leveling pads are raised to their highest position.**
- **Ensure that there is no stabilizer bracket installed on the rack cabinet.**
- **Do not use a ramp inclined at more than ten degrees.**
- **Once the rack cabinet is in the new location, do the following:**
  - **Lower the four leveling pads.**
  - **Install stabilizer brackets on the rack cabinet.**
  - **If you removed any devices from the rack cabinet, repopulate the rack cabinet from the lowest position to the highest position.**

If a long distance relocation is required, restore the rack cabinet to the configuration of the rack cabinet as you received it. Pack the rack cabinet in the original packaging material, or equivalent. Also, lower the leveling pads to raise the casters off of the pallet and bolt the rack cabinet to the pallet.

Handling static-sensitive devices

Attention: Static electricity can damage electronic devices, including your server. To avoid damage, keep static-sensitive devices in their static-protective packages until you are ready to install them.

To reduce the possibility of damage from electrostatic discharge, observe the following precautions:

- Limit your movement. movement can cause static electricity to build up around you.

- Handle the device carefully, holding it by its edges or its frame.
- Do not touch solder joints, pins, or exposed circuitry.
- Do not leave the device where others can handle and damage it.
- While the device is still in its static-protective package, touch it to an unpainted metal part of the server for at least 2 seconds. This drains static electricity from the package and from your body.
- Remove the device from its package and install it directly into the server without setting down the device. If it is necessary to set down the device, put it back into its static-protective package. Do not place the device on your sever cover or on a metal surface.
- Take additional care when handling devices during cold weather. Heating reduces indoor humidity and increases static electricity.

Notices and statements used in this document

The caution and danger statements that appear in this document are also in the multilingual *Safety Information* document, which is provided on the Web as a PDF document and an HTML document. Each statement is numbered for reference to the corresponding statement in the *Safety Information* document.

The following notices and statements are used in the documentation:

- **Note:** These notices provide important tips, guidance, or advice.
- **Important:**These notices provide information or advice that might help you avoid inconvenient or problem situations.
- **Attention:**These notices indicate possible damage to programs, devices, or data. An attention notice is placed just before the instruction or situation in which damage could occur.
- **Caution:**These statements indicate situations that can be potentially hazardous to you. A caution statement is placed just before the description of a potentially hazardous procedure step or situation.
- **Danger:**These statements indicate situations that can be potentially lethal or extremely hazardous to you. A danger statement is placed just before the description of a potentially lethal or extremely hazardous procedure step or situation.

Chapter 1. System overview

This chapter provides information about the operating systems that support the IBM Cluster 1350 components and related documentation.

The Cluster 1350 supports a maximum of 512 nodes in addition to the one required xSeries® 345 management node (or an @server 325 for 64-bit processing environments). All nodes must run one of the Linux versions shown in the following table.

Table 1. Linux operating-system support

Microprocessor	Operating-system support
32-bit Professional	Red Hat Linux version 9.0 (CSM), Red Hat Enterprise Linux (RHEL) version 3.0 (CSM)
32-bit Enterprise	SLES version 8 (CSM) for Opteron, Red Hat Enterprise Linux (RHEL) version 3.0 (CSM)
64-bit	SLES version 8 for Opteron, Red Hat Enterprise Linux (RHEL) for Opteron version 3.0 (CSM only), and Workstation for Opteron version 3.0 (XCAT)

The Cluster 1350 uses a primary cabinet and an expansion cabinet. The primary cabinet contains the management node and console monitor. An expansion cabinet can contain the following components:

- Cluster or compute nodes
- Storage nodes
- Mass storage devices

Note: An expansion cabinet does not contain a management node or console.

Figure 1 on page 2 illustrates a primary cabinet. Figure 2 on page 3 illustrates an expansion cabinet containing cluster nodes. Figure 3 on page 4 illustrates an expansion cabinet containing storage controllers and mass storage.

© Copyright IBM Corp. 2004

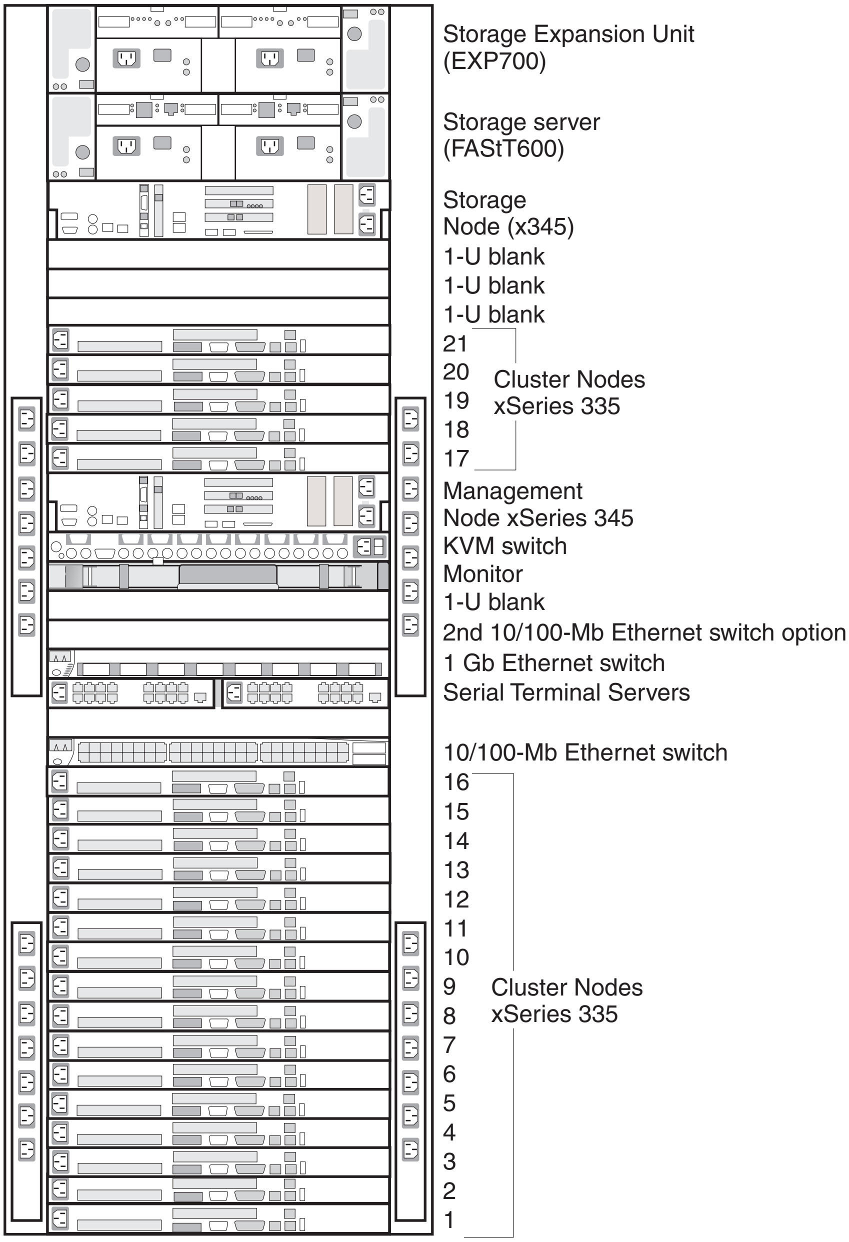
Storage Expansion Unit (EXP700)
Storage server (FAStT600)
Storage Node (x345)
1-U blank
1-U blank
1-U blank
21
20
19
18
17
Cluster Nodes xSeries 335
Management Node xSeries 345
KVM switch
Monitor
1-U blank
2nd 10/100-Mb Ethernet switch option
1 Gb Ethernet switch
Serial Terminal Servers
10/100-Mb Ethernet switch
16
15
14
13
12
11
10
9
8
7
6
5
4
3
2
1
Cluster Nodes xSeries 335

Figure 1. Example of an eserver Cluster 1350 primary cabinet

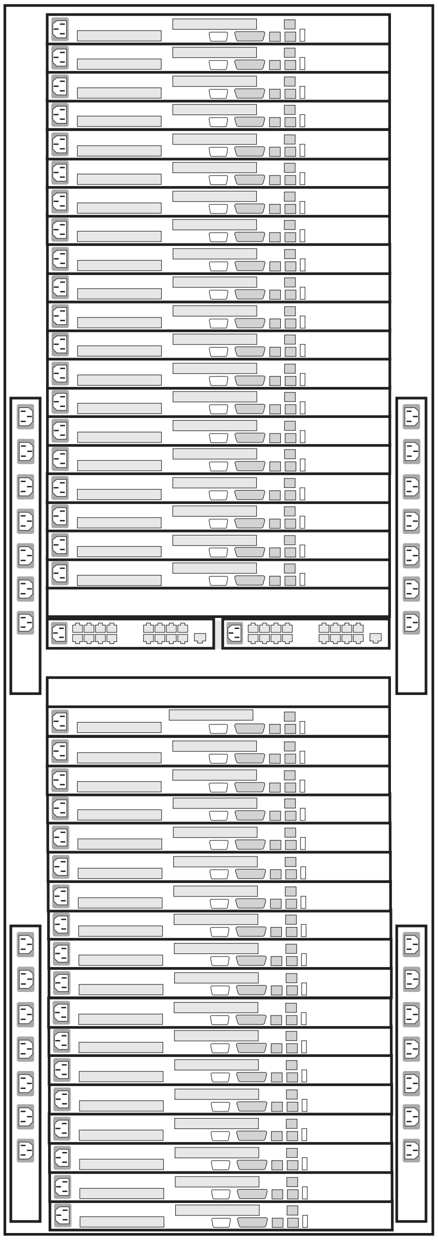

Cabinet 1	Cabinet 2	Cabinet 3	Cabinet 7	
38	76	114	blank	Cluster nodes xSeries 335
37	75	113	blank	
36	74	112	blank	
35	73	111	blank	
34	72	110	blank	
33	71	109	blank	
32	70	108	blank	
31	69	107	blank	
30	68	106	blank	
29	67	105	blank	
28	66	104	256	
27	65	103	255	
26	64	102	254	
25	63	101	253	
24	62	100	252	
23	61	99	251	
22	60	98	250	
21	59	97	249	
20	58	96	248	
19	57	95	247	
18	56	94	246	Cluster Nodes xSeries 335
17	55	93	245	
16	54	92	244	
15	53	91	243	
14	52	90	242	
13	51	89	241	
12	50	88	240	
11	49	87	239	
10	48	86	238	
9	47	85	237	
8	46	84	236	
7	45	83	235	
6	44	82	234	
5	43	81	233	
4	42	80	232	
3	41	79	231	
2	40	78	230	
1	39	77	229	

Figure 2. Example of an @server Cluster 1350 expansion cabinet with cluster nodes. This figure also shows how the node numbering scheme maps to other expansion cabinets.

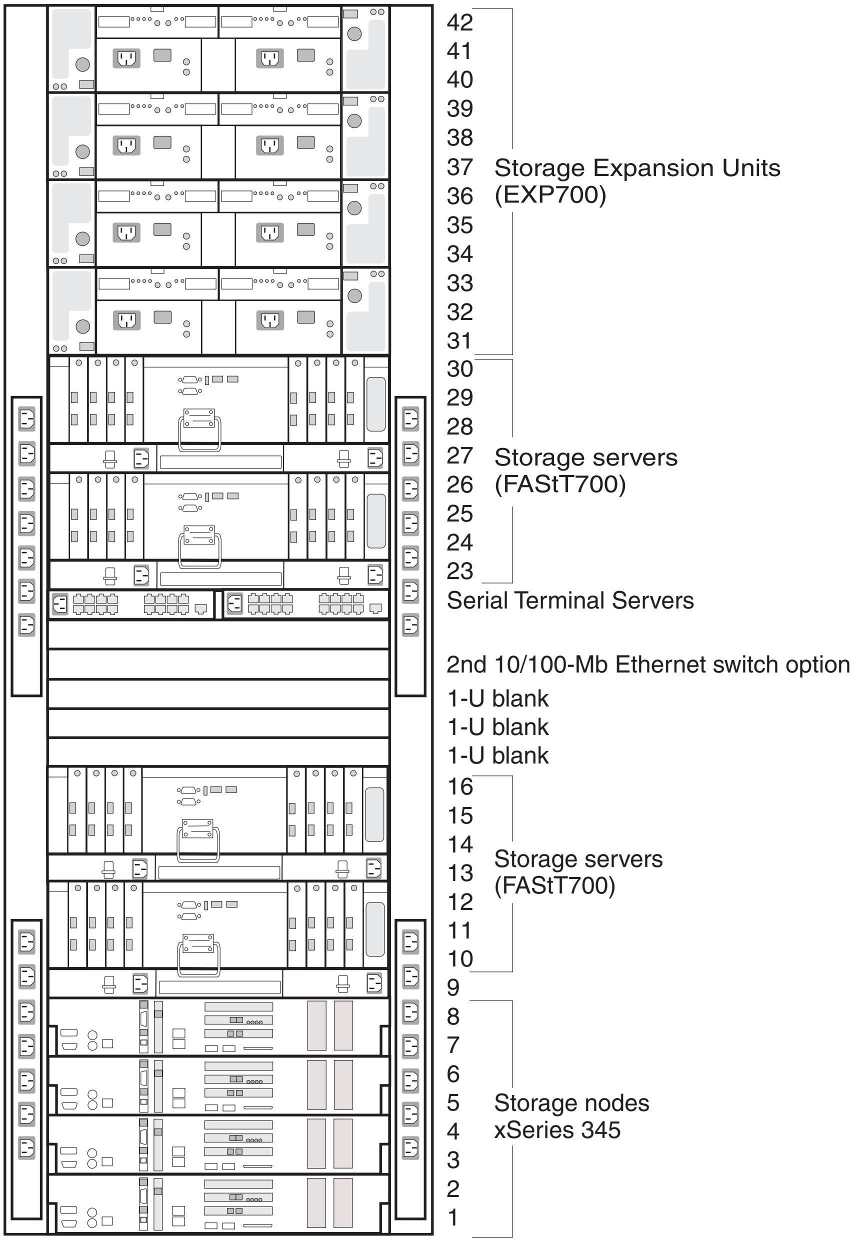
42
41
40
39
38
37
36
35
34
33
32
31
Storage Expansion Units (EXP700)
30
29
28
27
26
25
24
23
Storage servers (FAStT700)
Serial Terminal Servers
2nd 10/100-Mb Ethernet switch option
1-U blank
1-U blank
1-U blank
16
15
14
13
12
11
10
Storage servers (FAStT700)
9
8
7
6
5
4
3
2
1
Storage nodes xSeries 345

Figure 3. Example of an IBM @server Cluster 1350 expansion cabinet containing storage controllers and mass storage

Cluster components

This section describes the components in the Cluster 1350. Cluster components include:

- Cluster nodes
- Management node
- Storage nodes
- Storage servers
- Storage expansion units
- Serial ATA hard disk drive

- SCSI/RAID storage controller adapters
- Console
- KVM switch
- TopSpin InfiniBand switch
- Topspin InfiniBand host channel adapter
- 10/100 Ethernet switches
- 10/100/1000 Ethernet switch
- SMC mini-GBIC expansion module
- Terminal server
- High-speed Myrinet switches
- Myrinet BladeCenter connectivity
- Power distribution unit (PDU)

Cluster nodes

A cluster must contain at least four cluster nodes. The cluster nodes perform the computational tasks in the cluster. Cluster nodes are also known as compute nodes.

The cluster nodes are a combination of the following components:

- @server® 325
- xSeries 335
- xSeries 345
- @server BladeCenter™ unit populated with HS20 blade servers

See Table 1 on page 1 for more information about the Linux software versions supported on cluster nodes.

Management node

Each cluster contains one management node, which provides system management for all modules in the cluster. The Cluster 1350 management node is typically an xSeries 345 server running Linux. You can also use an @server 325 server as the management node in a cluster environment running a 64-bit Linux operating system.

Storage nodes

The optional storage nodes manage the mass storage. The cluster supports up to 32 storage nodes. The total number of storage and compute nodes cannot exceed 512.

For tasks that do not require large amounts of mass storage, the storage node has onboard disk storage that is typically sufficient. The storage nodes can be any of the following servers in the primary cabinet running the Linux operating system:

- @server 325 - a 1 U storage unit
- xSeries335 - a 1 U storage unit
- xSeries 345 - a 2 U storage unit
- xSeries 360 - a 3 U storage unit
- @server BladeCenter unit populated with HS20 blade servers

Storage servers

For the storage server option, the Cluster 1350 system communicates over a Fibre-Channel connection and uses any of the following RAID-capable controllers:

- FAStT200 storage server (3 U). Each FAStT200 storage server adds up to 10 internal 18 GB 15 000 RPM drives or ten 36 GB or 73 GB 10 000 RPM drives to the storage capacity of the cluster.
- FAStT600 storage server (3 U). Each FAStT600 storage server supports up to 14 internal disk drive modules, supporting over 2 TB of storage capacity when using 146 GB drives. Additional storage can be added to the FAStT600 with up to seven FAStT EXP700 storage expansion units using optional EXP700 attachment features.
- FAStT600 with Turbo storage server (3 U). Each FAStT600 with Turbo storage server supports over sixteen TB of fibre-channel disk using 7 EXP700s. You can configure a RAID protected storage solution to help provide up to 64 storage partitions.
- FAStT700 storage server (3 U). Each FAStT700 storage server supports up to 224 18 GB 15 000 RPM drives or 224 36 GB or 73 GB 10 000 RPM drives contained in external expansion cabinets.
- FAStT900 storage server (4 U). Each FAStT900 storage server supports up to 16 EXP700 external expansion cabinets. All servers interface with the cluster so that the storage nodes communicate with large RAID-protected arrays of storage.

Storage expansion units

The cluster supports the following disk storage expansion units:

- FAStT EXP100: supports up to 14 FAStT 250 GB SATA hard disk drives
- FAStT EXP400: supports up to 14 Ultra320 SCSI hard disk drives
- FAStT EXP500: supports up to 10 hard disk drives
- FAStT EXP700: supports up to 224 Fibre Channel hard disk drives
- FAStT EXP900: supports up to 220 Fibre Channel hard disk drives

Serial ATA hard disk drive

The IBM FAStT 250 GB dual-port SATA hard disk drive offers up to 3.5 TB of storage capacity per enclosure. Combined with FAStT EXP100 storage expansion units, you can configure RAID-protected storage solutions of up to 14 hard-disk drive modules.

SCSI/RAID storage controller adapters

The Cluster 1350 supports the following SCSI/RAID storage controller adapters:

- A ServeRAID™-6I Ultra320 SCSI controller supports up to 16 arrays with support for a maximum of 160 hard disk drives.
- A ServeRAID-6M Ultra320 SCSI controller supports eight arrays with support for a maximum of 30 hard disk drives.

Console

The console provides the monitor, keyboard, and mouse for the management node. The monitor is a 1 U flat-panel display that folds down and retracts into the rack.

KVM switch

The keyboard/video/mouse (KVM) switch allows the console to connect to all the nodes in the cluster from one terminal location. Storage and management nodes are connected directly to the KVM switch. For cluster nodes in the same rack, you can configure multiple nodes on one KVM switch port.

The Cluster 1350 can use one of the following devices:

- IBM NetBAY™ 2x8 console switch
- NetBAY Remote Console Manager (RCM)
- NetBAY Local Console Manager (LCM)

The RCM is the only supported KVM switch option that can be used with the *e*server 325 server.

TopSpin InfiniBand switch

The Topspin 120 InfiniBand switch provides either 24 InfiniBand ports of 10 Gbps connectivity or 8 InfiniBand ports of 30 Gbps connectivity in a single, 1-U unit.

Topspin InfiniBand host channel adapter

The Topspin InfiniBand Host Channel Adapter connects the Topspin 120 InfiniBand switch to the cluster enabling the InfiniBand switch to perform in a server cluster environment.

10/100 Ethernet switch

The 10/100 Mb Ethernet switch provides 10/100 Ethernet connections for the cluster. The Cluster 1350 uses the following 10/100 Ethernet switches:

- Cisco Catalyst Ethernet switch 3550 XL (24-port)
- Cisco Catalyst Ethernet switch 3550 XL (48-port)

You can partition the switch to set up multiple independent LANs within the same switch.

Each model also provides two 1 Gb Ethernet ports for communication with the management node.

10/100/1000 Ethernet switch

The 10/100/1000 Ethernet switch provides a 1 Gigabit Ethernet trunk line between the management node and the cluster and storage nodes. The 1 Gigabit uplink ports use optical cables. The Cluster 1350 uses the following switches as 10/100/1000 Mb Ethernet switches:

- IBM *e*server BladeCenter Gigabit Ethernet switch module (4-port)
- Cisco Catalyst Ethernet switch model 3750G-24T (24-port) stackable Gigabit Ethernet switch
- SMC Ethernet switch model 8624T(24-port)

SMC mini-GBIC expansion module

The SMC mini-GBIC expansion module provides a network connector that allows you to plug into any small form-factor (SPF) slot for Gigabit Ethernet network expansion.

Terminal server

The terminal server provides serial connections for cluster modules. The Cluster 1350 uses the following switches as terminal servers:

- iTouch IR-8020-101 (20-port) switch

- iTouch IR-8040-101 (40-port) switch
- In-Reach LX-4032 (32-port) switch
- In-Reach LX-4048 (48-port) switch
- Equinox CCM4850 console manager

The main purpose of the terminal server is to have out-of-band console access to cluster components.

High-speed Myrinet switches

This is an optional 2 Gb switch for interconnecting cluster nodes and storage nodes. The Cluster 1350 supports the following Myrinet models:

- M3-E32 (5-slot)
- M3-E64 (9-slot)
- M3-E128 (17-slot)
- M3F-PC164C-2 (PCI adapter)
- M3F-PCIXD-2 (low-profile PCI-X card)
- M3F2-PCIXE-2 (low-profile PCI-X dual-port card)

The high-speed switch can replace the optional secondary Ethernet switch. It requires a Myrinet PCI adapter in each cluster node and storage node. The Myrinet switch uses an optical cable.

Myrinet BladeCenter connectivity

Myrinet BladeCenter connectivity includes a Myrinet PCI card, the Optical Passthrough Module, and special cable options.

High-speed Cisco switches and line cards

The Cluster 1350 can use the Cisco Catalyst 4003 (3-slot) switch and the Cisco catalyst 4006 (6-slot) switch for a lower cost high-speed solution.

The Cisco Catalyst 6500 series switches and line cards deliver scalable performance and port density across a range of rack-unit configurations and LAN/WAN/MAN interfaces. The Cluster 1350 supports the following 6500 switches:

- 6503 - 3-slot cabinet that utilizes the 6548, a 48-port 10/100/1000 Mb Ethernet line card
- 6509 - 9-slot cabinet that utilizes the 6748, a 48-port fabric-enabled 10/100/1000 Mb Ethernet line card
- 6704 - 4-port 10 Gigabit Ethernet switch

Power distribution unit (PDU)

Each rack contains one or more of the following power distribution units:

- IBM NetBAY rack power distribution unit (PDU)
- IBM NetBAY front-end power distribution units
- IBM Distributed Power Interconnect (DPI®) rack power distribution units
- IBM Distributed Power Interconnect (DPI) front-end power distribution units
- IBM Distributed Power Interconnect (DPI) high density power distribution unit

Some PDUs are mounted sideways beside the rack space. The cluster supports rack PDUs and front-end PDUs. The cluster can have the PDUs installed in the side pockets and not accessible from the rear. Some of the newer cluster configurations provide access to the PDUs from the rear of the system cabinet.

Rack PDUs provide power to components within a cabinet; front-end PDUs provide the connection to the external power source and distribute the power among the rack PDUs. To eliminate the need for the front-end PDU, a rack PDU is directly connected to the external power source. Up to four front-end PDUs and up to 12 rack PDUs can be placed in each cabinet.

DANGER

> **The breaker switch on the PDU is not accessible. To turn off power to the cabinet, you must disconnect all the PDU power cords from the electrical outlets or from the individual PDU inlets.**

Related publications

Your cluster might have features that are not described in the documentation that you received with the cluster. The documentation might be updated occasionally to include information about those features, or technical updates might be available to provide additional information that is not included in your cluster documentation. These updates are available from the IBM Web site. Complete the following steps to check for updated documentation, related documentation, and technical updates:

1. Go to **http://www.ibm.com/pc/support/**.
2. In the **Learn** section, click **Online publications**.
3. On the “Online publications” page, in the **Brand** field, select **Servers**.
4. In the **Family** field, select **Clustering**.
5. Click **Continue** and select the online documents for the product.

Chapter 2. Unpacking the Cluster 1350

An IBM support team installs the IBM @server Cluster 1350. This chapter provides unpacking information about the 1410-42L (42 U), 1410-42X, and 1410-25X (25 U) rack cabinets and associated cluster components. These racks and the applicable components are installed by an IBM support team. The 1417-11X (11 U) rack cabinet and its components are installed by the customer.

Note: For more information about installing an IBM NetBAY11 rack (1417-11X) unit, see the *IBM NetBAY11 Rack Installation and Maintenance Guide* that came with your unit or go to http://www.ibm.com/pc/support, enter the product number in the Quick Path field, and click **Go**. The model and serial number are generally found on the back panel and invoice.

Complete the following steps before the IBM support team arrives on site to finish the installation:

1. Review the legal and safety information.
2. Review the physical, environmental, and electrical requirements in the *IBM @server Cluster 1350 Preinstallation Planning Guide* and make sure that the installation site is ready.
3. Unpack the cabinets only but not the other boxes. Depending on the cluster configuration that you ordered, the heavy cluster components are removed from the rack cabinets and packed separately to satisfy shipping requirements.

   Attention: Do not attempt to replace any equipment that was removed from the racks. The IBM support team will install all equipment back into its locations as part of the installation process.
4. Using the order that you placed for the cluster, identify the primary cabinet and verify its contents. If equipment is removed prior to shipping, check the bill of materials to make sure that all the equipment that is required for the primary cabinet was shipped with the order.
5. Using the order that you placed for the cluster, identify the expansion cabinets and verify their contents. If equipment was removed prior to shipping, check the bill of materials to make sure that all the equipment that is required for the expansion cabinets was shipped with the order.
6. Dispose of the packing material that comes with the cabinets.
7. Move the cabinets and any boxes containing extra equipment or other material to the installation site. The IBM support team completes the final cabinet placement.
8. Arrange for a phone line near the cabinets.

Note: The IBM support team performs the installation of components shipped outside of the cluster, the final cabling, and installation steps. After the IBM support team installs the cluster, connect the network cables.

© Copyright IBM Corp. 2004

Chapter 3. Placing the cabinets

This chapter provides information about the placement of the cluster cabinets and how to install the frame stabilizer foot and outrigger to support each cabinet.

Physical, environmental, and electrical requirements are outlined in the *IBM @server Cluster 1350 Preinstallation Planning Guide*. Do not move the cabinets to the installation location until the location is ready.

Customer responsibilities

After you have verified the contents of all the cabinets of the Cluster 1350, move the cabinets and any boxes containing extra equipment and other materials to the location that you have prepared for the installation.

Note: Your customer responsibilities include placement and installation of the 1410-11X (11 U) rack cabinet and its cluster components.

The Cluster 1350 is manufactured according to the information that you provided at the time you placed the order. The side-to-side and front-to-back clearances for each cabinet are directly related to the load-carrying capability of the floor in the installation location. Cabinet-to-cabinet cabling harnesses are custom made for each order according to the planned spacing of the cabinets. If the location of the installation changes from the location at the order time, review the physical, environmental, and electrical requirements outlined in the *Preinstallation Planning* manual to make sure that there are no incompatibilities.

Using the information in the packing slip enclosed with the Cluster 1350, place the cabinets in their approximate final locations. Each cabinet has installation labels to help you in this process.

The IBM support team determines the final cabinet placement and completes the cabling and installation steps.

Installer responsibilities

The instructions in this section are intended only for the IBM support team who will complete the installation of the 1410-42L, 1410-42X, and 1410-25x rack units and cluster components.

Use the following guidelines when placing the cabinets:

- Cabinets can be placed side-by-side in contact with one another. Remember that to service any power distribution unit (PDU) in a cabinet, you must remove the side covers. At least 30 inches of *working clearance* is required to ensure the safe removal of a side cover and provide access to the PDU. If the cabinets are placed side-by-side in contact with each other, leave enough extra space around the cluster so that you can move the cabinets if a PDU needs service. Cabinet placement must not exceed floor-loading limits.
- Cabinet placement must allow for access to both the front and back panels. At least 36 inches of *working clearance* is needed to remove or insert a module into the rack.
- Cables and cable harnesses are custom made to fit the order. If the location of the installation has changed since the time the order was placed, review the

© Copyright IBM Corp. 2004

physical, environmental, and electrical requirements outlined in the *Preinstallation Planning* manual to make sure that there are no incompatibilities.

- Make sure that the cabinets are arranged correctly and adjust them if necessary. See the packing slip and the cabinet labels to verify that all cabinets are in their correct locations.

DANGER

Ensure that all rack-mounted units are fastened in the rack frame. Do not extend or exchange any rack-mounted units when the stabilizer is not installed.

To finish the cabinet placement, complete the following steps:

1. Inspect the cabinets, components, and cable connections for shipping damage.
2. Install the frame-stabilizer foot on each cabinet. **Figure 4** shows how to install a frame-stabilizer foot.
3. Engage the rack side stabilizers on each cabinet. See the documentation that came with your rack.

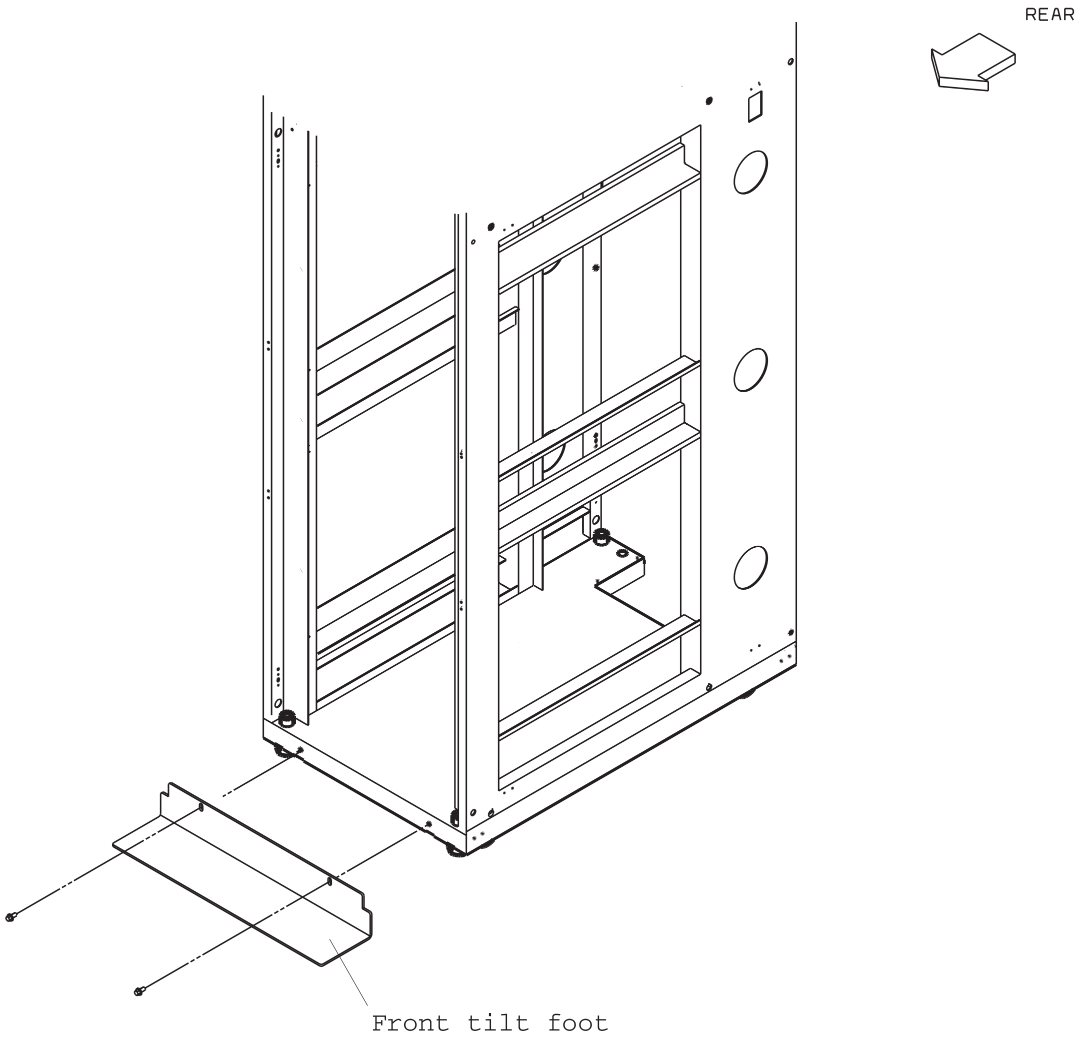

Figure 4. Installing the stabilizer (tilt) foot

Chapter 4. Cabling the Cluster 1350

Most of the cabling in a Cluster 1350 system is installed during manufacturing. However, there are three instances where cables must be installed at a customer site:

- Cables between cabinets
- Replacements for faulty cables
- Cables to replacement components

Any cable that fails at the customer site or is connected to components that must be replaced must be reconnected at the customer site.

The various types of cables in the Cluster 1350 system perform functions such as providing serial and Ethernet connections to cluster components.

Notes:

1. There are additional color-coded intercabinet Ethernet cables available to help you organize your cluster cabling by color. The current cable colors include green, blue, and yellow with lengths varying from 0.6m to 25m. These cables do not replace previous Ethernet cables but can be used in place of previous cables if you prefer a color-coded cabling scheme. Contact your sales representative to order additional color-coded intercabinet Ethernet cables.
2. In some clusters, the FRU interconnect cables can also be grey or white.

Management VLAN

The management VLAN provides the private virtual LAN (VLAN) to manage the components in the cluster. This VLAN includes the following connections:

- RS-485 connections to all cluster nodes and storage nodes through the Remote Supervisor Adapters. These enable diagnostics and monitoring for the cluster and storage nodes.
- Serial connections to all cluster components. These provide a path for configuration of components in the cluster.
- 10/100 Ethernet connections from the Remote Supervisor Adapter in the management node to the 10/100 Ethernet switch.

Primary cluster VLAN

The primary cluster VLAN provides a 10/100 or 10/100/1000 Ethernet connection (depending on the selected VLAN type) for communication with cluster nodes and storage nodes. This VLAN includes the following connections:

- A 10/100 or 10/100/1000 Mb Ethernet connection to all cluster and storage nodes and other components. This provides the primary communications between the management node and the other components in the cluster.
- A Gigabit Ethernet trunk line (shared with the management VLAN) for certain VLAN types only. This serves as a high-speed trunk line for all Ethernet communication within the cluster.

Optional secondary cluster VLAN

The optional secondary cluster VLAN provides a second 10/100/1000 Ethernet or a 2 Gb Myrinet switch for communication with cluster and storage nodes. The following options are available for the secondary cluster VLAN:

- A 10/100/1000 Ethernet connection

© Copyright IBM Corp. 2004

- A 2 Gb Myrinet connection

Keyboard/video/monitor

The keyboard/video/mouse (KVM) connects the ports on all nodes (cluster, storage, and management) to a single console through a central switch.

Fibre Channel cables

Fibre Channel cables provide Fibre-Channel connections between the storage nodes and the storage servers and between the storage servers and the storage expansion units.

Power distribution units

The power distribution unit provides the power to the cluster components. This includes both the power to the entire cabinet through the PDUs and remote power to the Remote Supervisor Adapter (RSA) boards and the terminal servers through the power management module.

VLAN options

The Cluster 1350 supports a variety of VLAN options. There are six basic configurations. Point-to-point wiring information is printed on each cable. Check the information on the cables in the primary rack and see the following tables to determine which VLAN option was used in the cluster.

Table 2. Type 1 10/100 Ethernet VLAN

Device	Management VLAN	10/100 primary cluster VLAN	Comments
Management node	Ethernet 2 connects to Cisco 3550		
KVM switch	Connects to Cisco 3550		
In-Reach LX-4000 terminal server (32-port, 48-port)	Connects to Cisco 3550		
APC switch	Connects to Cisco 3550		
Cisco 3550 10/100 switch			
Cluster nodes		Ethernet 0 connects to Cisco 3550	
Storage nodes	Ethernet 2 connects to Cisco 3550		
FAStT600	Connects to Cisco 3550		Uses both jacks
FAStT700	Connects to Cisco 3550		Uses both jacks
FASt900	Connects to Cisco 3550		Uses both jacks

Table 3. Type 2 10/100/1000 Ethernet VLAN

Device	Management VLAN	Gbit primary cluster VLAN	Comments

Table 3. Type 2 10/100/1000 Ethernet VLAN (continued)

Management node	• Ethernet 1 connects to Cisco 400x • Ethernet 2 connects to Supervisor I or Supervisor III		06P3701 or 22P7801
KVM switch	Connects to Cisco 400x		
In-Reach LX-4000 (32-port, 48-port) terminal server	Connects to Cisco 400x		
APC switch	Connects to Cisco 400x		
Cisco 4003 switch, 4006 switch, or both	• Supervisor I connects to management-node • Supervisor III connects to management-node	• Gbit connects to management-node Ethernet1 • Supervisor III uplink connects to management-node PCI adapter	
Cluster nodes		Ethernet 0 connects to Cisco 400x	
Storage nodes	Ethernet 2 connects to Cisco 400x	Ethernet 1 connects to Cisco 400x	
FAStT600	Connects to Cisco 400x		Uses both jacks
FAStT700	Connects to Cisco 400x		Uses both jacks
FAStT900	Connects to Cisco 400x		Uses both jacks

Table 4. Type 3 VLAN: 10/100 Ethernet with 10/100/1000 public high-speed VLAN

Device	**Management VLAN**	**10/100 primary cluster VLAN**	**Gbit customer public high-speed VLAN**
Management node	Ethernet 2 connects to Cisco 3550		
KVM switch	Connects to Cisco 3550		
In-Reach LX-4000 terminal server (32-port, 48-port)	Connects to Cisco 3550		
APC switch	Connects to Cisco 3550		
Cisco 4003 and/or 4006 switch	• Supervisor I connects to Cisco 3550 • Supervisor III connects to Cisco 3550		
Cluster nodes		Ethernet 0 connects to Cisco 3550	Ethernet 2 connects to Cisco 400x Gbit
Storage nodes			Ethernet 2 connects to Cisco 400x Gbit
FAStT600	Connects to Cisco 3550		
FAStT700	Connects to Cisco 3550		
FAStT900	Connects to Cisco 3550		

Table 5. Type 4 VLAN: 10/100/1000 Ethernet with 2 Gbit public high-speed VLAN

Device	**Management VLAN**	**10/100 primary cluster VLAN**	**Myrinet customer public high speed VLAN**
Management node	Ethernet 2 connects to Cisco 3550		
KVM switch	Connects to Cisco 3550		
In-Reach LX-4000 terminal server (32-port, 48-port)	Connects to Cisco 3550		
APC switch	Connects to Cisco 3550		
Myrinet 32, 64, or 128 (both jacks), D card	Connects to Cisco 3550		
Topspin InfiniBand switch	Connects to Cisco 3550		
Cluster nodes		Ethernet 0 connects to Cisco 3550	Myrinet adapter
Storage nodes	Ethernet 2 connects to Cisco 3550		Myrinet adapter
FAStT600	Connects to Cisco 3550		
FAStT700	Connects to Cisco 3550		
FAStT900	Connects to Cisco 3550		

Table 6. Type 5 VLAN: 10/100/1000 Ethernet with 10/100/1000 Ethernet public high-speed VLAN

Device	**Management VLAN**	**Gbit primary cluster VLAN**	**Gbit customer public high-speed VLAN**
Management node	• Ethernet 1 Alias connects to Cisco 400x • Ethernet 2 connects to Supervisor I or Supervisor III	• Ethernet 1 connects to 4003 • Fibre Channel PCI adapter connects to 4006 Supervisor III uplink	Copper or Fibre Channel PCI connects to public network
KVM switch	Connects to Cisco 400x		
In-Reach LX-4032 and 4048 (32-port, 48-port) terminal server	Connects to Cisco 400x		
APC switch	Connects to Cisco 400x		
Cisco 4003 switch, 4006 switch, or both	• Supervisor I connects to management-node Ethernet 2 • Supervisor III connects to management-node Ethernet 2	• Gbit connects to management-node Ethernet 1 • Supervisor III uplink 1 connects to Fibre Channel PCI adapter	Supervisor III uplink 2 connects to public network

Table 6. Type 5 VLAN: 10/100/1000 Ethernet with 10/100/1000 Ethernet public high-speed VLAN (continued)

Cluster nodes		Ethernet 0 connects to Cisco 400x	Ethernet 2 connects to Cisco 400x
Storage nodes	Ethernet 1 Alias connects to Cisco 400x	Ethernet 1 Alias connects to Cisco 400x	Ethernet 2 connects to Cisco 400x
FAStT700 (both jacks)	Connects to Cisco 400x		

Table 7. Type 6 VLAN: 10/100/1000 Ethernet with 2 Gbit Myrinet public high-speed VLAN

Device	**Management VLAN**	**Gbit primary cluster VLAN**	**Myrinet customer public high-speed VLAN**
Management node	• Ethernet 1 Alias connects to Cisco 400x • Ethernet 2 connects to Supervisor I or Supervisor III	• Ethernet 1 Alias connects to Cisco 400x • Fibre Channel PCI adapter (part number 06P3701 or 22P78021) connects to Cisco 4006 Supervisor III uplink	
KVM switch	Connects to Cisco 400x		
In-Reach LX-4032 and 4048 terminal server (32-port, 48-port)	Connects to Cisco 400x		
APC switch	Connects to Cisco 400x		
Myrinet 32, 64, 128, or PCI adapter (both jacks)	Connects to Cisco 400x		
Topspin InfiniBand switch	Connects to Cisco 400x		
Cluster nodes		Ethernet 0 connects to Cisco 400x	Myrinet adapter
Storage nodes	Ethernet 2 connects to Cisco 400x	Ethernet 1 connects to Cisco 400x	Myrinet adapter
FAStT700	Connects to Cisco 400x		

For large clusters that use VLAN type 2, 5, or 6, you must install additional Cisco 400x switches. If the cluster has more than one Cisco 400x switch in the primary rack, refer to the following tables.

Table 8. Type 2 VLAN with multiple Cisco 400x switches

Device	**Management VLAN**	**Gbit primary cluster VLAN**	**Comments**
Management node	Ethernet 2 connects to Cisco 3550		
KVM switch	Connects to Cisco 3550		

Table 8. Type 2 VLAN with multiple Cisco 400x switches (continued)

In-Reach LX-4032 and 4048 (32-port, 48-port) terminal server	Connects to Cisco 3550		
APC switch	Connects to Cisco 3550		
Cisco 3550 10/100 switch		Cisco 3550 copper uplink	
Cisco 4003 switch, 4006 switch, or both	• Supervisor I connects to 3550 • Supervisor III connects to 3550		
Cluster nodes		Ethernet 0 connects to Cisco 400x	
Storage nodes	Ethernet 2 connects to Cisco 3550	Ethernet 1 connects to Cisco 400x	
FAStT700	Connects to Cisco 3550		Uses both jacks

Table 9. Type 5 VLAN with multiple Cisco 400x switches

Device	Management VLAN	Gbit primary cluster VLAN	Gbit customer public high-speed VLAN
Management node	Ethernet 2 connects to Cisco 3550		Fibre Channel PCI connects to public network
KVM switch	Connects to Cisco 3550		
In-Reach LX-4032 and 4048 (32-port, 48-port) terminal server	Connects to Cisco 3550		
APC switch	Connects to Cisco 3550		
Cisco 3550 10/100 switch	Connects to management- node Ethernet 2		
Cisco 4003 switch, 4006 switch, or both	• Supervisor I connects to 3550 • Supervisor III connects to 3550		
Cluster nodes		Ethernet 0 connects to Cisco 400x	Ethernet 2 connects to Cisco 400x
Storage nodes	Ethernet 1 Alias connects to Cisco 400x	Ethernet 1 Alias connects to Cisco 400x	Ethernet 2 connects to Cisco 400x
FAStT700 (both jacks)	Connects to Cisco 3550		

Table 10. Type 6 VLAN with multiple Cisco 400x switches

Device	Management VLAN	Gbit primary cluster VLAN	Myrinet customer public high-speed VLAN

Table 10. Type 6 VLAN with multiple Cisco 400x switches (continued)

Management node	Ethernet 2 connects to Cisco 3550		Fibre Channel PCI adapter connects to public network
KVM switch	Connects to Cisco 3550		
In-Reach LX-4032 and 4048 (32-port, 48-port) terminal server	Connects to Cisco 3550		
APC switch	Connects to Cisco 3550		
Cisco 3550 10/100 switch	Connects to management-node Ethernet 2		
Cisco 4003 switch, 4006 switch, or both	• Supervisor I connects to Cisco 3550 • Supervisor III connects to Cisco 3550		
Myrinet 32, 64, 128, or PCI adapter (both jacks)	Connects to Cisco 3550		
Topspin InfiniBand switch	Connects to Cisco 3550		
Cluster nodes		Ethernet 0 connects to Cisco 400x	Myrinet adapter
Storage nodes	Ethernet 2 connects to Cisco 3550	Ethernet 1 connects to Cisco 400x	Myrinet adapter
FAStT700	Connects to Cisco 3550		

Connecting the cables

Cables and the cable harnesses in each cabinet are labeled with information that tells where to connect each end of the cable. Each label identifies the device or node it connects to, and where applicable, its port number.

Depending on the country of manufacture the label scheme will vary. Before you begin attaching cables, become familiar with the information on the labels.

When installing a Cluster 1350, start with the primary cabinet. After you have connected the intracabinet cables inside the primary cabinet, move on to each expansion cabinet and use the information printed on each cable label to connect the cables in the cabinet.

After you have connected any cables in the primary cabinet and expansion cabinets, connect the cables that run between the cabinets. This is called the intercabinet cabling, and the following types of cables are involved:

- 1 Gb or 2 Gb Fibre Channel (optical)
- 1 Gb Ethernet (optical)
- 2 Gb Myrinet (optical)
- 10/100/1000 Ethernet (copper)
- KVM (copper)
- Remote console manager (RJ-45/Cat 5)

- Local console manager (Cat5E)
- Topspin 120 InfiniBand server (RS232)
- Topspin InfiniBand host channel adapter

For a complete listing of all available cables and their part numbers see the Cluster 1350 information on the IBM InfoTips Web site: http://w3.pc.ibm.com/helpcenter/infotips/. The information is also available in the IBM Current Object Repository (CORE) system.

The following sections describe the different types of cabling.

1 Gigabit Ethernet cabling

The 1 Gb Ethernet provides a high-speed optical trunk line for VLAN communication with the management node.

Each intercabinet cable has labels at both ends. You can use the information on the label to create a site map to document all cable routing.

Myrinet high-speed switch cabling

The Myrinet high-speed switch provides an optional 2 Gb optical network for communications between cluster nodes and storage nodes.

Each intercabinet cable has labels at both ends. You can use the information on the label to create a site map to document all cable routing.

10/100/1000 Ethernet cabling

The 10/100/1000 Ethernet switch provides an optional 10/100/1000 network for communications between cluster nodes and storage nodes.

Each intercabinet cables have labels at both ends of each cable. You can use the information on the label to create a site map to document all cable routing.

Fibre Channel cabling

Fibre Channel is used to connect storage nodes to storage servers and to connect storage servers to storage expansion units.

Each intercabinet cable has labels at both ends. You can use the information on the label to create a site map to document all cable routing.

KVM cabling

The KVM switch allows a maximum of eight connections. Use the following guidelines for cabling the KVM switch:

- Use the information on each end of each cable to create a site map.
- When routing a KVM cable from a cabinet containing cluster nodes to another cabinet containing the KVM switch, connect a C2T-to-KVM cable to the cluster nodes and use a KVM extension cable to add sufficient length to reach the KVM switch.
- Multiple KVM switches can be connected in series. Cluster nodes (xSeries 335 server) can be connected in series to a single KVM switch port (up to 40 cluster nodes), but the management node and all the storage nodes each require a separate KVM switch port. Certain systems might require a second KVM switch. Install the second switch in the expansion cabinet that contains the additional storage nodes.

- When using two KVM switches, connect Port A (the console port) of switch 2 to port 8 of switch 1. Use a KVM extension cable to make the connection between the two cabinets. If more length is needed, use two KVM extension cables linked together.

Remote console manager cabling

The remote console manager (RCM) switch has 16 ACT connections (KVM over RJ-45/CAT5) and one KVM connection for the console. Use the following guidelines for cabling the RCM switch:

- Use the information on each end of each cable to create a site map.
- When routing a Cat 5 KVM cable from a cabinet containing cluster nodes to the cabinet containing the RCM, use a CCO cable and a Cat 5 cable sufficiently long enough to reach the RCM switch.
- Multiple KVM switches can be connected in series.
- Up to 40 cluster nodes (xSeries 335 server) can be connected in series to each ACT port on the RCM. The management node and all the storage nodes can also be daisy-chained, with up to 16 per ACT port. Multiple RCMs can not be daisy-chained together. The RCM can be connected to an Ethernet network to allow for remote access to the consoles of the servers over the network.

Local console manager cabling

The local console manager (LCM) switch accepts the CAT5 input from the NetBay Conversion Options and allows them to be fed into your management station. Its four CAT5 input ports support four chains of 16 servers per chain. The local console manager (LCM) switch also supports Cable Chaining Technology (C2T).

Use the following guidelines for cabling the LCM switch:

- Use the information on each end of each cable to create a site map.
- When routing a Cat5 KVM cable from a cabinet containing cluster nodes to the cabinet containing the LCM, use a CCO cable and a Cat5 cable sufficiently long enough to reach the LCM switch.
- Multiple KVM switches can be connected in series.
- Up to 16 cluster nodes (xSeries 335 server) can be connected in series.

Topspin 120 InfiniBand switch and host channel adapter cabling

Each Topspin 120 InfiniBand switch accepts one RS-232 cable and one copper or optical cable for the Ethernet management port. The cables can be ordered in 3 meter or 10 meter lengths.

Replacing a defective cable in a harness

If a cable in a harness is defective, complete the following steps to replace the cable:

1. Disconnect both ends of the defective cable from their ports. Do not remove any other connectors from their ports.
2. If possible, remove the cable from the harness. Otherwise, use a pair of wire cutters to cut off the connectors at both ends of the defective cable. This prevents someone from mistakenly reconnecting the cable, thinking that it has inadvertently been left unconnected.
3. Install a single cable between the two empty ports. Use a wire tie to attach the cable to the harness that contains the defective cable. This identifies the replacement cable as belonging to this harness.
4. Label the replacement cable so it is clearly identified as a replacement.

Chapter 5. Checking the Cluster 1350 cabling process

The IBM @server Cluster 1350 comes without an operating system installed. An IBM support team performs the hardware cabling process then the customer performs the software installation unless a service contract is purchased that includes the software installation.

Before turning on a Cluster 1350 system, you must first, check all the connections in the expansion cabinets and primary cabinet. After you have verified that all connections are secure, turn on the expansion cabinets containing storage nodes, storage servers, and storage expansion units. Turn on the primary cabinet last.

Checking connections in the expansion cabinets

1. Make sure that the breaker switches for the source power are all turned off.
2. Open the side and rear doors of the cabinet.
3. From the side of the cabinet, make sure that all the power cables between the rack power distribution units and the front-end power distribution units (PDUs) are fully seated.
4. From the back of the cabinet, push on all the connectors on the cables running from the rack-mounted devices powered by the power distribution units to make sure that the cables are fully seated.
5. Connect power to the power distribution units:
   a. Connect the power cable to the power distribution unit.
   b. Draw the power cable through the opening at the base of the cabinet.
   c. Connect the power cable to the electrical outlet.
   d. Turn on the power breaker switch for the source power.
   e. Make sure that the power distribution unit circuit breakers are turned to the **On** position.
6. Make sure that all internal power distribution units are turned on by viewing the power LEDs on the power distribution unit connected components.
   - When power is applied, servers display a flashing green LED on the front panel.
   - The following devices have no power switch and turn on automatically when the power distribution units are turned on. Make sure that the following components have power applied:
     – KVM switches
     – Cisco 10/100 switches
     – Cisco Gigabit switches
     – In-Reach LX-4032 and LX-4048 (32-port, 48-port) terminal servers

   All rack-mounted devices are powered by the internal power distribution unit.

Checking connections in the primary cabinet

Complete the following steps to check the connections in the primary cabinet:

1. Make sure that the source power breaker switches are all turned off.
2. Open the side and rear doors of the cabinet.
3. From the side of the cabinet, make sure that all the power cables between the rack power distribution units (PDUs) and the front-end PDUs are fully seated.
4. From the back of the cabinet, push on all the cable connectors running from the rack-mounted devices powered by the power distribution units to make sure that the cables are fully seated.
5. Connect power to the power distribution units. Use a NEMA L6-20, 280 V ac, single-phase power cable.

© Copyright IBM Corp. 2004

a. Connect the power cable to the power distribution unit.
b. Draw the power cable through the opening at the base of the cabinet.
c. Connect the power cable to the electrical outlet.
d. Turn on the power breaker switch for the source power.
e. Make sure that the power distribution unit circuit breakers are turned to the **On** position.

6. Make sure that all internal power distribution units are turned on by viewing the power-on LEDs on the power distribution unit connected components.
   - When power is applied, servers display a flashing green LED on the front panel.
   - The following devices have no power switch and turn on automatically when the power distribution units are turned on. Make sure that the following components have power applied:
     - KVM switch
     - Cisco 10/100 switch
     - Cisco Gigabit switch
     - In-Reach LX-4000 (32-port, 48-port) terminal server

   All rack-mounted devices are powered by the internal PDU.

Turning on the power to the expansion cabinets

Complete the following steps to turn on the expansion cabinets:

1. Turn on the cluster nodes using their power switches.
2. Make sure that every Remote Supervisor Adapter has power. A green LED on the board is lit to indicate that the adapter has power.
3. After an expansion cabinet is turned on, make sure that all front panel LEDs on the cluster nodes are lit, otherwise, not all the nodes will be shown in the configuration.

Repeat the procedure for every expansion cabinet unit in the cluster before powering on the primary cabinet.

Turning on the power to the primary cabinet

Complete the following steps to turn on the primary cabinet:

1. Storage expansion units - on the back of each storage expansion unit turn on the circuit breakers.
2. Storage controllers - on the back of each storage controller turn on the circuit breakers.
3. On the front of each management node, turn on the power switch. Make sure that the node passes POST with no errors. During the startup process make sure that the PXE boot agent utility program attempts to run. If it does not, press F1 to start the Configuration/Setup utility program and add **Network** as a third boot option. After the node has started, the **No operating system** icon is displayed correct. If any yellow warning LEDs on the management node are lit, fix the underlying condition before continuing.
4. Turn on cluster node 1 and make sure that the node passes POST with no errors. During the startup process make sure that the PXE boot agent utility program attempts to run. If it does not, press F1 to start the Configuration/Setup utility program and add **Network** as a third boot option. After the node has started, the **No operating system** icon is displayed correct. If any yellow warning LEDs on the management node are lit, fix the underlying condition before continuing.
5. On the front of each storage node, turn on the power switch. During the startup process make sure that the PXE boot agent utility program attempts to run. If it

does not, press F1 to start the Configuration/Setup utility program and add **Network** as a third boot option. After the node has started, the **No operating system** icon is displayed correct. If any yellow warning LEDs on the management node are lit, fix the underlying condition before continuing. You must turn on the peripheral devices and bring them online before you turn on the storage nodes so that the storage nodes can detect them.

6. Make sure that all Cat 5 and fibre channel connections have a green link LED.
7. Make sure that each Remote Supervisor Adapter has power. A green LED on the adapter faceplate is lit to indicate that power is applied to the adapter. Connect your mobile computer to the Cisco 10/100 switch and configure it to use IP address 172.22.30.20 with a net mask of 255.255.0.0. Log in to each Remote Supervisor Adapter using the Web interface and make sure that each adapter is present. Ping each communication device (Cisco switches, In-Reach LX-4000 (32-port, 48-port) terminal server, power management module, and KVM switch).

Verifying the installation of the Linux Cluster Installation Tool

A startable CD is included with your Linux Cluster Installation Tool (LCIT) installation materials. With your 1350 Cluster, you also receive tab files that list the rack hardware configuration. To make sure that you install the cluster components correctly, run LCIT to generate a new set of tab files. Compare the new tab files to the tab files that come with your cluster to make sure that the two sets of tab files are identical. If the tab files are not identical recheck your component installation. You can also use LCIT to help you locate Remote Supervisor Adapters or nodes that do not start as expected. Complete the following steps to verify the installation of LCIT:

To run the Linux cluster installation tool, complete the following steps:

1. In the rear of the rack, unplug all compute nodes.
2. Turn on the management node and insert the LCIT CD into the CD-ROM drive. Workspace 1 opens as a gray screen.
3. Anywhere in the open Workspace 1 window, right-click and select **LCIT 3.2**. The LCIT window opens.
4. Right-click on the IP address **172.20.0.1**, and then select **Alias** to set the Baseboard Management Controllers. The factory default is 172.29.0.1 with 16-bit subnet.
5. Click **Configuration** and click **Apply** to write the dhcpd.conf file.
6. Plug in the power cords for all the compute nodes in chronological order. Node numbering increases from the bottom of the rack upward and from left to right in the BladeCenter unit.
7. Start each server and make sure that the servers display the management-node Workspace 1 gray screen. This step validates that each server has loaded the ram-disk image from the network.

   Note: If you do not see a Workspace 1 gray screen on a server, make sure that the server is configured to boot from the network.

Lights out or brown out event

The following sequence occurs during a lights out or brown-out event scenario.

1. The system is up and running typical applications.
2. A lights out or brown out event occurs. The system turns off and then turns back on through an external source.
3. All nodes turn on to the last known state (On/Off). If the last known state is On, then the nodes start and display a login prompt.
4. Log files show system restart events on nodes and on Remote Supervisor Adapter. If a lights out or brown out event occurs, check the following log files:
   - /var/log/messages
   - /var/log/csm/installnode.log (management server)
   - /var/log/csm/install.log (on the node)
   - Remote Supervisor Adapter event log
   - BIOS event log

Related topics

See Appendix C, “Error and event logs,” on page 91.

Chapter 6. Installing the software

These installation procedures are intended for IBM Global Services or the customer's agent for the initial software setup of a Cluster 1350. Complete the following steps before proceeding:

1. Install a supported version of Linux. See Cluster 1350 Drivers, Firmware, and Software Levels at http://publib.boulder.ibm.com/cluster/current.htm for all supported versions of Linux and other supported versions of Cluster 1350 software and firmware.
2. Install Cluster Systems Management (CSM) software.
3. Configure the storage nodes.
4. Distribute the system image to all nodes in the cluster.
5. Test the configuration.

The installation time is approximately 8 hours per cabinet.

Before you begin the software installation process, see "Software version matrix" to verify that you have all the required material.

Important: The Cluster 1350 must be maintained only by system administrators who are experienced with Red Hat Enterprise Linux, DHCP, NFS, and Linux networking and administration.

Software version matrix

The Cluster 1350 requires certain levels of a supported Linux version and Cluster System Management (CSM) software. Before you begin the software installation process, make sure that you have collected all the applicable levels of operating-system kernel, management software, device drivers, and other firmware that are needed for building a working system image. This up-to-date information is at http://publib.boulder.ibm.com/cluster/index.html.

If the versions of any components are changed, IBM might not be able to service and support the cluster.

Downloading device drivers and firmware

If you need device drivers and firmware, go to http://www-3.ibm.com/pc/support/site.wss/multiplefiledownload.do, select your machine Brand, Family, Type and Model, if applicable, and click **Continue** to locate your drivers.

Installing a supported version of Linux

The Cluster 1350 comes without an operating system. The customer or the customer's agent is responsible for securing a valid copy of the operating system for installation. See Cluster 1350 Drivers, Firmware, and Software Levels at http://publib.boulder.ibm.com/cluster/current.htm for more information about the supported versions of Linux.

Use the detailed installation instructions that come with your software kit to install the Linux software. If you do not have your documentation for installing Linux, go to: http://www.redhat.com/docs/manuals/linux/.

© Copyright IBM Corp. 2004

Installation of the operating system begins with the management node in the primary cabinet.

Installing the Cluster System Management software

To install the Cluster System Management (CSM), see the installation instructions that come with your CSM software kit. You can also obtain installation information: http://www-1.ibm.com/servers/eserver/clusters/library/am7LXstp.pdf.

Installing the General Parallel File System system-management software

To install the General Parallel File System (GPFS) software, see the installation instructions that come with your GPFS software kit. You can also obtain installation information: http://www-1.ibm.com/servers/eserver/clusters/library/gpfs.html.

Configuring the storage nodes

The procedure assumes that the following prerequisites have been met:

- A supported version of Linux is installed and running from a local drive. See "Software version matrix" on page 31.
- The FAStT storage server is properly configured and connected to a host bus adapter.
- The FAStT drives in the storage server are configured into different RAID groups, storage groups, and LUN through the software that comes with the FAStT storage server.

Because of the way the Red Hat Enterprise Linux version 3.0 loads SCSI drivers and assigns them to dev/sda and dev/sdb, problems can result if more than one SCSI host adapter board (Adaptec SCSI controller for local drives or QLogic for Triton connection) is installed on the system and you use the `scsi_hostadapter` alias. When the system is restarted, the operating system detects the QLogic controller before detecting the Adaptec controller, which will cause the system problems. To avoid this problem, follow the procedure in Installing the storage node software and modify the order of the contents of the etc/modules.conf file.

Installing the storage node software

Complete the following steps to install storage node software:

1. If you have not already done so, turn off the storage controllers or disconnect the Fibre-Channel cable running to each storage controller.
2. For Red Hat Enterprise Linux version 3.0, edit the etc/modules.conf file to put the host adapters in the correct order and to add the parameter `scsi_mod max_scsi_luns` to the file.

   Important:

   - Because the system is running a modular kernel, the Adaptec SCSI device driver (alias scsi_hostadapter aic7xxx) must be probed before any other SCSI adapters so that the kernel will be able to find the root device during the initialization phase. Also, the QLogic QLas device driver (alias scsi_hostadapter2 qla2x00) must be the last SCSI host adapter listed in the file.
   - If there are other SCSI host adapter boards installed on the system and the scsi_hostadapter alias is used, define a different alias for the qlogic Qla driver and be sure to add it after the other SCSI modules so that the SCSI devices names already in use are not renumbered the next time the system is

started. You can do this by appending a number at the end of the scsi_host adapter word, for example, alias scsi_hostadapter*n* qla2x00 (where *n* is an alphanumeric number from 1 through 9).

The original module.conf file might be similar to the following example:

```
alias eth0 e1000
alias scsi_hostadapter qla2x00
alias scsi_hostadapter1 aic7xxx
alias scsi_hostadapter2 ips
alias eth1 e1000
alias parport_lowlevel parport_pc
alias scsi_hostadapter3 aic7xxx
alias scsi_hostadapter4 aic7xxx
alias usb-controller usb-ohci
```

The modified module.conf file might be similar to the following example:

```
alias eth0 e1000
alias scsi_hostadapter1 aic7xxx
alias scsi_hostadapter2 ips
alias eth1 e1000
alias parport_lowlevel partport_pc
alias scsi_hostadapter3 aic7xxx
alias scsi_hostadapter4 aic7xxx
alias scsi_hostadapter5 qla2x00
alias usb-controller usb-ohci
options scsi_mod max_scsi_luns=128
```

For SLES8: edit the /etc/modules.conf file so that it contains the following lines:

```
alias scsi_hostadapter ips
alias scsi_hostadapter1 qla2300
options scsi_max_scsi_luns=128
```

For SLES8, edit the /etc/rc.config file to contain the following line:

```
INITRD_MODULES="ips qla2300"
```

3. For Red Hat Enterprise Linux version 3.0, rebuild the two initrd images. You can not make a RAM disk image if mkinitrd detects one already present with the same name, so the first two commands will rename the old images. For example:

```
mv /boot/initrd-2.4.2-2.img /boot/initrd-2.4.2-2_orig.img
mv /boot/initrd-2.4.2-2smp.img /boot/initrd-2.4.2-2smp_orig.img
mkinitrd initrd-2.4.2-2.img 2.4.2-2
mkinitrd initrd-2.4.2-2smp.img 2.4.2-2smp
```

   For SLES8, type the mkinitrd command to create a boot/initrd directory, and then run lilo.

4. If a Remote Supervisor Adapter is installed, restart and run the setup diskette or CD to configure the network. Assign the same configuration information for the Remote Supervisor Adapter (name, IP address, host name) as used before. Go to the following site to download the RSA and ASM Process or Firmware Update Diskette utility: http://www.pc.ibm.com/qtechinfo/MIGR-4JTS2T.html
5. If you have custom modifications, configure the kernel.
6. Reboot the node.

Defining nodes

Before you can distribute the system image to all nodes in the Cluster 1350, you must first define the nodes in the /etc/hosts file, assigning each node an IP address. After you complete this step, you must define the nodes in CSM using the `definenode` command.

Defining nodes using CSM

You can use a node-definition file to define the nodes, console servers, and service processors to the cluster, or you can enter the information from the command line.

1. To define the nodes, console server information, and service processors, type **`definenode -f nodedef`** and press **Enter**.
2. To review the arguments that you need to enter from the command line, type **`definenode -h`** and press **Enter**.
3. To define the node host name, type **`definenode -n`** ***`hostname`*** and press **Enter**. Where *hostname* is the name of the node being defined. The command prompts for missing information when some or all of the arguments are not provided.

See the *IBM Cluster Systems Management for Linux Technical Reference* for details about definenode and addnode command-line syntax and more examples of the usage of the command.

Defining nodes using GPFS

Use the mmaddnode command to add nodes to an existing GPFS node set. On each new node, a mount point directory and character mode device is created for each GPFS file system.

Syntax: mmaddnode [-C NodesetId] {-n NodeFile | NodeName[:manager | client][,NodeName[:manager | client]...]}

Parameters:

__ `C NodesetId`

The identifier of the GPFS node set you want to add nodes to. If this option is not specified or a period (.) is used for the *NodesetId*, the nodes are added to the node set from which the mmaddnode command was issued. To determine the node set, at the command prompt, type

```
mmlsnode -C
```

__ `-n` *NodeFile*

The file containing the list of node descriptors, one per line, to be added to the node set. Node descriptors are of the same format as in the `NodeName[:manager | client]` parameter, and follow the same rules.

__ `NodeName[:manager | client][,NodeName[:manager | client]...]`

A comma-separated list of nodes to be added to the node set. Nodes are specified by a node name and can be optionally followed by a use designation. A designation of manager specifies that the node must be included in the pool of nodes from which the file system manager node is chosen. For further information on the role of a node as the file-system manager, see the *General Parallel File System for Linux(R): Concepts, Planning, and Installation Guide* and search for *file system manager*.

The host name or IP address must refer to the communications adapter. Alias interfaces are not allowed. Use the original address or a name that is resolved by the host command to that original address. You can specify a node using any of these forms: `Short hostname` *k145n01*
`Long hostname` *k145n01.kgn.ibm.com*
`IP address` *9.119.19.102*

Security

You must have root authority to run the mmaddnode command. You can issue the mmaddnode command from any node in the GPFS cluster.

When using rcp and rsh for remote communication, a properly configured rhosts file must exist in the root user's home directory, normally /root, on each node in the GPFS cluster. If you have designated the use of a different remote communication program in the mmcrcluster or the mmchcluster command, make sure that proper authorization is granted to all nodes in the GPFS cluster and the nodes in the GPFS cluster can communicate without the use of a password.

Examples

1. To add nodes k145n04 and k145n05, designating k145n04 to be available as a manager node only, and by default, add the nodes to the GPFS node set, type:
   ```
   mmaddnode k145n04:manager,k145n05
   ```
2. To confirm the addition, type: `mmlsnode -C`.
3. To add nodes k145n06 and k145n07 to the GPFS node set set1, type:
   ```
   mmaddnode -C set1 k145n06,k145n07
   ```
4. To confirm the addition, type `mmlsnode -C set1.`

Rules to follow when adding nodes

You must follow these rules when adding nodes to a GPFS node set:

- A node can belong to only one node set at a time.
- The nodes being added to the node set must belong to a GPFS cluster (issue the mmlscluster command to display available nodes).
- The existing nodeset must meet quorum for the nodes to be added. For example, if GPFS is currently configured on eight nodes, all of which are up and running, the quorum value is met and the new nodes join the node set.
- Conversely, if GPFS is currently configured on eight nodes and only four are up and running, a quorum of five does not exists and the new nodes can not join the node set. When five of the original eight nodes are up and running, the new nodes are added.
- After the nodes have been added and GPFS is started on the new nodes, the quorum value for the nodeset is adjusted accordingly. This enables new nodes to join a running node set without causing quorum to be lost.
- Issue the mmstartup command to start GPFS on the new nodes.
- When adding nodes to a node set using the single-node quorum algorithm, the GPFS daemon must be stopped on all of the nodes. If after the nodes are added, the number of nodes in the node set exceeds two, the quorum algorithm is automatically changed to the multinode quorum algorithm.

Distributing the system image to all nodes in the cluster

Because of the way the Red Hat version 9.0 loads SCSI drivers and assigns them to dev/sda, dev/sdb partitions, problems can result if more than one SCSI host adapter (Adaptec or LSI SCSI controller for local drives and QLogic HBA for Triton connection) is installed on the system. The QLogic HBA will typically be detected first by the installation process. Follow the "Installing the storage node software" on page 32 and modify the order of the contents of the /etc/modules.conf file.

Attempting to distribute the system image out to the nodes while a FAStT controller is still turned on and connected might cause data damage on the first logical disk device in FAStT subsystem. Make sure that the FAStT controllers are turned off or that all fiber cables for the FAStT controllers are disconnected from the back of each controller before starting the install process.

To distribute the system image to all nodes in the cluster, complete the following steps:

1. Open an rconsole window for each node being installed so you can monitor the installation process:

   ```
   rconsole -n {node_list}
   ```

2. Run the following setup command:
   - For SLES, type:**csmsetupyast**
   - For Red Hat, type:**csmsetupks**
3. Run the `installnode` command for each node being installed:

   ```
   installnode {node_list}
   ```

After the operating system is installed on the storage nodes, reconnect the fibre cable to the FAStT controllers. Restart the storage nodes to see any configured LANs.

Verifying the configuration

1. Start the management node and log on as user. At the command prompt, type: **`root`**.
2. Log on to the storage nodes and verify the disk configuration:
   ```
   fdisk -l
   ```
3. If a modem is present, configure the modem according to the instructions.

Chapter 7. Administering the Cluster 1350

This chapter includes information about:

- accessing the cluster from a remote location
- accessing each node before the operating system is installed
- shutting down the system components
- finding information about a lights out or brown out event

For more information about monitoring, remote control, set-up, and technical references, see: http://www.ibm.com/servers/eserver/clusters/library/linux.html

Using the remote power command

Using the command rpower starts and resets hardware, powers hardware on and off, and queries the node power state. The syntax is:

```
rpower -h
[-v] {-a | -A |
{[-n node_list | [-N nodegroups]
[-d device_list] [-D devicegroups]}}
{on | off | reboot | [-l] query | resetsp_hcp |
resetsp_host | -m {full | lpar} cec_on | cec_off | [-l] cec_query}
```

Remote console

The remote console function uses the serial ports of xSeries servers and terminal servers. The serial ports provide remote access to nodes before the operating system is installed or when network access to the servers is unavailable or has failed. Terminal servers must be included to enable the remote console function.

Each rack in the configuration includes one or two terminal servers to connect each node in the rack through a DB9 to RJ45 serial cable. The terminal servers are LAN connected to the Management VLAN.

The remote console function is accessed through the rconsole command. This command opens a remote console session for any cluster node. The syntax is:

```
rconsole [-h] [-x]
[-a | -A | {[-n node_list] [-N nodegroups]
[-d device_list] [-D devicegroups]}]
[-t] [-o] | -O number of columns]] [-c [-v] | [-r | -f]]
```

Shutting down the system components

Because the operating system is installed on each node, the procedures to shut down the nodes are relatively simple. Complete the following steps to shut down the cluster nodes:

1. Log off the cluster nodes and the storage nodes.
2. To turn the cluster nodes off with Cluster System Management (CSM) installed on the management node, at the command prompt, type

   ```
   rpower -a off
   ```

   Note: If CSM is not installed on the management node, turn off the power switch for each individual cluster node.

© Copyright IBM Corp. 2004

3. Turn off the power for the following devices:
   - storage nodes
   - management node
   - storage controllers
   - storage expansion units
4. Turn off the power switch for the power distribution units or unplug, from the power distribution unit, the devices that have no power switch.
   - To turn off the power distribution units, unplug the power cables from the wall outlet.
   - The following devices have no power switch and must be unplugged if the power distribution units are not turned off:
     – KVM switch
     – Cisco 10/100 switch
     – Cisco 10/100/1000 Mb switch
     – In-Reach terminal server
5. Unplug the power distribution unit power cables from the wall outlet.

Finding information about a lights out or brown out event

In the event of a lights out or brown out event, the following sequence occurs:

- The system is up and running typical applications.
- A lights out or brown out event occurs. The system shuts down then restarts through an external source.
- All nodes turn on to the last known state (On/Off). If the last known state is On, then the nodes will boot to a console login prompt.
- Log files show system restart events on nodes and on Remote Supervisor Adapter devices. Review the following log files for more information about the event:
  – /var/log/messages
  – /var/log/csm/installnode.log (management server command)
  – /var/log/csm/install.log (node command)
  – RSA event log
  – BIOS event log

Related topics

- Chapter 5, “Checking the Cluster 1350 cabling process,” on page 27
- Appendix C, “Error and event logs,” on page 91

Chapter 8. Troubleshooting hardware and software problems

This chapter includes information to diagnose problems associated with the Cluster 1350. The Cluster 1350 is an integrated Linux cluster that includes IBM and third party hardware and software components like server nodes and associated firmware, storage and networking subsystems, plus Cluster Systems Management (CSM) software and General Parallel File System (GPFS) software.

Problem resolution involves identifying the problem cluster component and following the applicable problem resolution procedure for that component.

This chapter includes information for the diagnosis of problems down to the component level. When a failing component is identified you can review the specific product documentation for further actions. Links to applicable product Web sites and online product documentation are provided in this chapter.

Diagnosing hardware and software problems in a cluster environment requires a basic understanding of how the components of the Cluster 1350 function together.

The cluster consists of:

- One or more racks.
- From 4 to 512 cluster nodes. The nodes are configured to execute customer applications or provide other services required by the customer, such as, file server, network gateway, or storage server.
- One management node (xSeries 345 or @server 325) for cluster systems management and administration.
- A management Ethernet VLAN used for secure traffic for hardware control.

  The management Ethernet VLAN is used for management traffic only. It is logically isolated for security using the VLAN capability of the Cisco Ethernet switches, and is only accessible from the management node. The cluster VLAN and management VLANs share the same physical Cisco switches.
- A cluster VLAN used for other management traffic and user traffic. Cisco switches integrated with the cluster are used for the management Ethernet VLAN and the cluster Ethernet VLAN.
- Service processor networks. All nodes in the cluster are connected through serial service processors (xSeries 335) and/or Remote Supervisor Adapter (RSA) devices. The first node in a serial connection must have a Remote Supervisor Adapter which is connected through the Ethernet to the management-Ethernet VLAN.
- A terminal server network for remote or local console. Optionally, the customer might elect to include an additional network.
- A high-performance Myrinet 2000 cluster interconnect, or an additional 10/100 Ethernet.
- The customer can elect to configure a subset of cluster nodes with additional external storage. This can also be a Fibre Channel solution (using a FAStT storage subsystem).
- A supported distribution of the Linux operating system.
- Cluster management software, such as, CSM.

CSM maintains a database of configuration information (tab files) about the nodes that are configured in the Cluster 1350. To display the node configuration information, use the following CSM command on the management server console:

```
lsnode -l
```

© Copyright IBM Corp. 2004

The output provides information about each node, such as, the node type, model number, serial number, and host name. The tab file output also provides information that corresponds each node to its terminal server network and service processor network. For the terminal server network, the output includes the console server host name and the console port number to which the node is connected. For the service processor network, the output includes the host name of the Remote Supervisor Adapter device to which the node is connected and the internal service processor name for the node.

To display a list of pre-managed nodes on the management server, at the console prompt, type

```
Isnode -a Mode
```

To display a list of managed nodes on the management server, at the console prompt, type

```
Isnode
```

To display the management server, at any node console prompt, type

```
mgmtsvr
```

Isolating network, node, and Linux problems

Cluster 1350 nodes are connected over a 10/100 Mb Ethernet cluster network. A Cluster 1350 can also have a second network, either an additional Ethernet network or a Myrinet 2000 network.

As a preliminary diagnostic step, ping all the nodes over all available networks.

Compare the error to possible symptoms in Table 11.

Table 11. Troubleshooting the shared VLAN

Symptom	Action
• Can ping the storage node from the management node but cannot ping the cluster nodes. • Can ping the cluster nodes from the management node but cannot ping the storage nodes. • Cannot ping either the storage nodes or the cluster nodes. • Cannot ping the cluster nodes in one of the expansion cabinets.	1. Make sure that the links between the management node, storage nodes, Cisco 3550, 3500, and 400x switches. 2. Restart the management node and press F1 to enter the Setup utility. Make sure that the Ethernet devices are turned on. 3. Make sure that the correct driver level is installed for 1Gb Fibre Channel Ethernet. To view the status, at the console prompt, type, `ifconfig`. 4. Check the Cisco (3500, 4003, or 4006) switch for green status LEDs or system and status LEDs. If the green LEDs are lit, the switch is OK. 5. To make sure that the 1 Gb Fibre-Channel Ethernet connections are good, swap a known good cable with a suspect cable to isolate the failing device. 6. Replace the failing Fibre-Channel cable, GBIC, or network interface card.

If following the steps in Table 11 did not correct the problem, continue with the steps shown in “Clustering with one network” on page 43.

Clustering with one network

Use this section for more troubleshooting information about clustering with one network.

Ping failure on one or some nodes

If one or more nodes experience a ping failure, this indicates a problem with the node hardware or software. Complete the following steps to resolve the problem:

1. Telnet or use (shell script) SSH to connect to the node through the serial console or KVM and make sure the node is operational.
   a. If telnet or SSH succeeds, check the syslog for errors.
      1) If there are errors, go to "Isolating software problems" on page 52 and complete the steps in that table to resolve the problem.
      2) If there are no errors, this indicates a network problem. See Table 12 and complete the steps in that table to resolve the problem.
   b. If telnet or SSH fails, this indicates a node hardware problem. See "Isolating hardware problems" on page 46 for problem resolution.

Ping failure on all nodes

If all nodes experience a ping failure, it indicates a problem on one of the following:

- Network. Go to Table 12 and complete the listed actions to resolve the problem.
- Network adapter on the management node
- DHCP configuration
- Network configuration
- Cisco blade failure

Table 12. Network troubleshooting for a cluster with one network

Symptom	Action
Cannot ping a node or nodes on the cluster network from the management node, yet the rconsole command and access from the KVM work correctly.	1. At the console prompt, type the `ifconfig` command to Make sure that the IP settings are correct. 2. Make sure that the cables are fully plugged into the switch and node, and that everything else is plugged into the correct port. See the cabling information printed on each cable label and "VLAN options" on page 18 if you are unsure where a cable belongs. Make sure that the link LEDs are lit. 3. Swap ports on the Ethernet switch with a known working cluster node port. 4. Make sure that the Ethernet switch port is configured for the Management VLAN.

Clustering with two networks

Use the following troubleshooting information about clustering with two networks.

Ping failure on one or more nodes

If one or more nodes experience a ping failure, it indicates a problem with the node hardware or software. Complete the following steps to resolve the problem:

- Telnet or use (shell script) SSH to connect to the suspect node via the serial console or KVM and make sure that the node is operational. If telnet or SSH succeeds, check the syslog for errors.
  1. If there are errors, go to "Isolating software problems" on page 52 and complete the listed actions to resolve the problem.
  2. If there are no errors, this indicates a network problem. Go to Table 14 on page 45 and complete the listed actions to resolve the problem.

- If telnet or SSH fails, this indicates a node hardware problem. Go to “Isolating hardware problems” on page 46 and complete the listed actions to resolve the problem.

Ping failure on only one network

If ping failures occur on one network but not on the other network, this indicates a problem on the network adapter on the management node for the failing network.

Ping failure on one or both networks

1. Make sure that all communication devices on the network are turned on and that each device has a green status LED lit on both ends of the connection.
2. Make sure that you have the correct IP Address, Net Mask, and Gateway settings for each device that fails to function in the network.
3. To determine the IP Address scheme of each node, at the console prompt, type, `ifconfig` and compare this output to the factory defaults shown in Table 13.

Table 13. Factory defaults

Device	IP address	Host name
Management node	• cluster VLAN: 172.20.0.1 • management VLAN: 172.30.0.1	eth0 – mgt.cluster.com eth1– mgt1.cluster.com
Management node alias	172.29.101.1	bmc
e325 node bmcs	172.29.0.1	bmc001
Storage node	172.20.1.1	storage001
FAStT storage controller	172.30.2.1	triton001
xSeries 335 cluster node	172.20.3.1	node001...node*xxx*
BladeCenter Ethernet switch module	172.30.101.1	sm001/mm001
Myrinet switch	172.30.10.1	myri001
Topspin InfiniBand switch	172.30.10.1	top-001
Terminal server	172.30.20.1	ts001
Remote Supervisor Adapter (bottom device)	172.30.30.1	rsa001
Remote Supervisor Adapter (next device)	172.30.30.2	rsa002
Remote Supervisor Adapter (Myrinet switch)	172.30.30.3	rsa003
Cisco 10/100 Mb switches	172.30.40.1	cisco3550-001
Cisco Gigabit Ethernet switches	172.30.50.1	• cisco3508-001 • cisco3750-001
SMC switch	172.30.50.1	smc8624-001
Cisco 6500 series switch	172.30.80.1	• cisco6503-001 • cisco6509-001
Cisco 4003 and 4006 switch (console management)	172.30.80.1	• cisco4003–001 • cisco4006–001
Remote console manager	172.30.70.1	rcm001

Ping failure on all nodes on both networks

If all nodes on both networks experience a ping failure, it indicates a problem with the system software or a user application. Telnet or use SSH to connect to the node through the serial console:

1. If telnet or SSH succeeds, check the *syslog* for errors.
   a. If there are errors, go to “Isolating software problems” on page 52 for software problem resolution.
   b. If there are no errors, this indicates a user application problem.
2. If telnet or SSH fails, connect to the node using a serial communications program, like, Hyperterminal. If you still cannot connect it indicates a node hardware problem. Go to “Isolating hardware problems” on page 46 for problem resolution.

Table 14. Network troubleshooting for a cluster with two networks

Symptom	Action
Cannot ping a node or nodes on the cluster network from the management node, yet the rconsole command and access from the KVM work correctly.	1. Use the ifconfig command to make sure that the IP address settings are correct. 2. Make sure that the cables are fully plugged into the switch and node, and that everything is plugged into the correct port. See the cabling information printed on each cable label and“VLAN options” on page 18if you are unsure where a cable belongs. Make sure that the link LEDs are lit. 3. Swap ports on the Ethernet switch with a known working cluster node port. 4. Make sure that the Ethernet switch port is configured for the management VLAN.

Isolating hardware problems

This section includes troubleshooting information about cluster hardware problems.

Node checks

Table 15. Troubleshooting the remote console network

Symptom	Action
1. Cannot execute a rconsole command to any cluster node. 2. Cannot execute any rconsole commands to get an active terminal session.	1. Make sure that the Ethernet connections between the terminal server and the Cisco switch are OK. Also check the connections between the Cisco switch and the management node. 2. Check the cables, dongles, and connectors at the node and the terminal server. Make sure that the serial port at the node is attached to the serial port on the terminal server by using the CSM command: `Isnode-aI <NodeName>`. See the ConsolePortNum information later in this section. 3. Follow steps 1 through 9 of Chapter 11, “Configuring a terminal server after device replacement,” on page 65, then at the `IN-Reach_Priv>` prompt, type, `show port <portnumber>` to compare the settings of all suspect ports against ports that are working. 4. Make sure that the terminal server is turned on and connected to the network by pinging the unit at the IP address of 172.30.20.1. 5. To make sure the serial port (COM 1) is configured to redirect output to the terminal server, complete the following steps: a. Restart the node and review the console screen. b. When the message, `Press F1 for Configuration/Setup` opens, press **F1**. c. From the main menu, select **Devices and I/O Ports** then press **Enter**. d. Make sure that Serial Port A is set to **Port 3F8, IRQ 4**. e. Select **Serial Port A**. f. Select **Remote Console Redirection**. g. Check the following settings: Remote Console Active [Enabled] Remote Console Com Port [COM1] Remote Console Baud Rate [9600] Remote Console Data Bits [8] Remote Console Parity [None] Remote Console Stop Bits [1] Remote Console Emulation [VT100] Remote Console Active After Boot [Enabled] h. Save the settings and exit. 6. For the xSeries 335 and xSeries 345 only, run diagnostics against the serial port to validate network connectivity. 7. Swap out the cables and dongle with new cables and dongle.

If the procedures in Table 15 on page 46 do not correct the issue you may have a problem with a port on the terminal server. Complete the following steps to test a different port:

1. Issue the CSM command `lsnode –I <nodename> lgrep Port` and record the port information.
2. Move the cable to a new port and change the port number using the CSM command `chnode <nodename> ConsolePortNum=xx` where *xx* is the new port number.

If the symptom persists, go to "Checking service processor logs" on page 54 and check the service processor log.

Hardware problem in service processor log

Go to "Node checks" on page 46 for node problem resolution.

Amber LED lit on node

The service processor log might be full. The log is cleared by connecting to the service processor through the Remote Supervisor Adapter card. Otherwise, go to "Node checks" on page 46 for node problem resolution.

`rpower` to node fails

Use the following information to resolve remote power failures:

- Check the service processor connection.
- At the console prompt, type `rpower -a on`.
- Go to "Checking service processor logs" on page 54 and check the service processor log.
- Use the Web interface, telnet, or SSH to reach each Remote Supervisor Adapter card on the cluster and then connect to the service processor on each cluster node individually through the Remote ASM Access menu.
- If the cluster node is not in listed, insert the node firmware diskette and try to diagnose the problem.

Service processor network

This section includes troubleshooting information about service processor network problems.

Table 16. Troubleshooting the service processor network

Symptom	Action
1. The `rpower –a query` command does not return with the status of all nodes 2. Cannot see all the nodes when managing remote Advanced System Management (ASM) service processors. 3. Cannot connect to individual Remote Supervisor Adapter (RSA) cards using browser.	1. Check the physical connections on the RS485 network and check for errors. 2. From the management node, use the Web browser and try to connect to the failing node through that node's RSA card. 3. Check that the RSA network is properly terminated. When more than one node is connected, terminators should be plugged into the empty port on the dongle and in the second RS485 port of the last node in the chain. 4. Swap the internodal Cat 5 cable on the unresponsive node with a known good cable. Also, replace the dongle if a problem is suspected. 5. Swap the KVM/RS485 cable (on the xSeries 335 only) with a known good cable. Also, replace the dongle if a problem is suspected. 6. Review the RSA configurations and IP settings with support. 7. Check the 10/100 Ethernet link between the RSA card and the Cisco 3550 or 400x switch. 8. Flash the ASM service processors to the latest firmware level. 9. Flash the RSA to the latest firmware level. 10. Check RSA configurations using the firmware update diskette.

If following the steps in Table 16 did not correct the problem, continue with the steps shown in "Remote Supervisor Adapter card connection failure."

Remote Supervisor Adapter card connection failure

Use the following information to resolve Remote Supervisor Adapter card connection failures:

1. To make sure the node has power, type: `rpower query`.
   a. If the node has power, `ping` the Remote Supervisor Adapter (RSA) card using the `HWControlPoint` field in the `lsnode` output.
      1) If `ping` succeeds, reset the RSA card, type: `rpower -n <nodename> resetsp_hcp`. If the adapter connection continues to fail after it has been reset contact IBM support.
      2) To check the logs, type: `reventlog -n <nodename> all`
      3) If the `ping` fails, check the network connection.
   b. If the node does not have power, check the power connections.
2. If the network connection LED is lit for the RSA adapter at the Ethernet switch, go to "Resetting the Remote Supervisor Adapter card" on page 54 and reset the RSA adapter.

Node connection or command failure

If the Remote Supervisor Adapter (RSA) card connection is working, but the node connection or commands issued to the node failed:

- Connect to the RSA card and check the node list.
- Check all cabling.
- Go to “Node checks” on page 46 and perform node checks.

Checking storage

Use the following information to resolve fibre-storage network failures:

Table 17. Troubleshooting the fibre-storage network

Symptom	Action
Cannot see disk drives from the storage node.	1. Reboot the storage node and press the **Alt/Q** keys to go into Qlogic setup. Make sure that the FAStT700 is a listed device. 2. Check the fibre connections between the server host bus adapter (HBA) and the hubs on the FAStT700. The green connection LED should be lit. 3. Check cabling on the FastT outbound hubs to the storage expansion units. Look for link lights and proper cabling. Also make sure that all transfer rate speed switches are set to 2 Gigabytes. 4. Check the Blade Server and Enhanced System Manager (ESM) firmware levels and update to current levels.

If following the steps in Table 17 did not correct the problem, continue with the steps shown in “File system failure.”

File system failure

Check disks using `fdisk -l`:

- If `fdisk -l` completes without error, go to “GPFS checks” on page 53 and continue with the file system problem resolution.
- If `fdisk -l` reports missing disks, check that the adapter device driver is configured:
  - If the adapter device driver is configured, go to “Checking storage” and continue with storage subsystem problem resolution.
  - If the adapter device driver is not configured, check the adapter hardware and then refer to the applicable documentation that came with your software and complete the problem resolution process.

PFA alert indicates internal disk

Go to “Checking storage” and perform disk problem resolution.

I/O errors in syslog

Complete the problem resolution for the indicated disk, adapter, or storage subsystem.

Ping failure over the Ethernet

Check the nodes using the `rconsole` command or `ping` nodes using the Myrinet switch:

- If the node responds, refer to the applicable documentation that came with the switch and complete the problem resolution process.
- If the node does not respond, go to “Node checks” on page 46 and continue with Node checks.

Ping failure over a Myrinet switch

Check the nodes using the `rconsole` command or `ping` nodes via the Ethernet:

- If node responds, go to “Configuring SNMP alerts from Myrinet” on page 53 and continue with Myrinet problem resolution.
- If the node does not respond, go to “Node checks” on page 46 and continue with Node checks.

Checking the terminal server

Check the terminal server nodes using the telnet SSH command or ping the nodes using the Ethernet:

- If node communication fails, go to “Node checks” on page 46 and continue with Node checks.
- If following the steps in Table 18 do not correct the problem, continue with the steps shown in “Isolating network, node, and Linux problems” on page 42..

Table 18. Troubleshooting the terminal server network for the Remote Console

Symptom	Action
Unable to run `rconsole` commands to get an active terminal session.	1. Check connection of cables and connectors at the nodes and the In-Reach terminal server. 2. Make sure that the In-Reach terminal server is turned on and connected to the network by pinging the unit at 172.30.20.1. 3. Follow steps 1 through 9 of Chapter 11, “Configuring a terminal server after device replacement,” on page 65, then at the `*IN-Reach_Priv>*` prompt, type, *`show port`* <*`portnumber`*> to compare the settings of all suspect ports against ports that work. 4. Make sure that the In-Reach terminal server is powered up and functional by pinging the unit at 172.30.20.1 5. Make sure that the serial port (COM 1) is configured to redirect the output to the terminal server: • Restart the node and watch the monitor screen. • When indicated, press **F1**. • From the main menu, select **Devices and I/O Ports**, then press **Enter**. • Make sure that Serial Port A is set to **Port 3F8, IRQ 4**. • Select **Serial Port A**. • Select **Remote Console Redirection**. • Check the following settings: Remote Console Active [Enabled] Remote Console Com Port [COM1] Remote Console Baud Rate [9600] Remote Console Data Bits [8] Remote Console Parity [None] Remote Console Stop Bits [1] Remote Console Emulation [VT100] Remote Console Active After Boot [Enabled] • Save settings and exit. 6. Swap out cables and dongle with known good units. 7. Run diagnostics against the serial port to test the connectivity.

Troubleshooting the KVM network

Use the following troubleshooting information to resolve KVM network failures:

Table 19. Troubleshooting the KVM network

Symptom	Action
The Keyboard/Video/Mouse (KVM) switch selector shows some or all systems are non-active (indicated by a red X) but the system is turned on.	1. Check that the connections for the KVM harness on the back of the system are securely plugged in. 2. Check the connection of the inbound/outbound Cat 5 connections on the KVM switch conversion dongle. 3. Check that the link LED on the dongle is lit. If the LED is lit, a good connection exists with the node keyboard port. If no link LED is lit and you are having problems with KVM connectivity, replace the dongle (FRU 32P1654). 4. Make sure that the terminator is in place at the first dongle on the KVM chain. 5. Use a known good Cat 5 (straight through) cable to direct connect or bypass possible bad cables. 6. Reboot the failing node to reset connection to the KVM switch.

File-system failure

Use the following information to resolve file-system failures:

To check the disks, at the command prompt, type: `fdisk -l`

- If `fdisk -l` completes, go to “GPFS checks” on page 53 and continue with the file system problem resolution.
- If `fdisk -l` reports missing disks, check that the adapter device driver is configured:
  - If the adapter device driver is configured, go to “Checking storage” on page 49 and continue with storage subsystem problem resolution.
  - If the adapter device driver is not configured, check the adapter hardware and then go to the applicable Linux documentation that came with your software and resolve the configuration issue.

PFA alert indicates internal disk

Go to “Checking storage” on page 49 and perform disk problem resolution.

I/O errors in syslog

Complete the problem resolution for the indicated disk, adapter or storage subsystem.

Isolating software problems

Use the following information to isolate and resolve software problems.

Operating-system checks

Use the following information to resolve operating-system checks:

Node non-responsive

If the node does not respond to `ping` or the serial console, and there are no relevant entries in the `syslog` or hardware logs, refer to the applicable Linux documentation that came with your software to continue with the problem resolution process.

Adapter device driver not configured

If the device driver is not configured, and there are no adapter hardware problems reported, refer to the applicable Linux documentation that came with your software and continue with the problem resolution process.

CSM checks

Use the following information to resolve CSM checks:

Events not logged or actions not taken

Using the ERRM command line interface, monitor the `AnyNodeProcessorsIdleTime` condition on specific managed nodes with the `LogEventsAnyTime` response while causing arm and rearm events. If arm and rearm events are not observed at the management server, this is configuration or network problem. Refer to the applicable documentation that came with CSM and complete the problem resolution process.

Differences in node lists

Output from the command `CT_CONTACT=<ManagedNodeName> lsrsrc IBM.[Host|FileSystem]` when run on the management node is not the same as when run on the managed node. This is configuration or network problem, refer to the applicable documentation that came with CSM and complete the problem resolution process.

netstat output incomplete

The command `netstat -an | grep rmc` on the management server does not show *ESTABLISHED TCP* connections for each managed node that is currently turned on. This is configuration or network problem, refer to the applicable documentation that came with CSM and complete the problem resolution process.

RMC not running

The command `lssrc -ls ctrmc` shows that RMC is not running on the management server. Refer to the applicable documentation that came with CSM and complete the problem resolution process.

lsrsrc reports errors

The command `lsrsrc -ab IBM.[Host|FileSystem]'` which checks that HostRM and FSRM will run on the management server reports errors. Refer to the applicable documentation that came with CSM and complete the problem resolution process.

lsaudrec reports errors

The command `lsaudrec` which checks that AuditRM will run on the management server reports errors. Refer to the applicable documentation that came with CSM and complete the problem resolution process.

Predefined conditions not shown

The `lscondition` and `lsresponse` commands when run on the management server do not show pre-defined conditions and responses. Refer to the applicable documentation that came with CSM and complete the problem resolution process.

Commands or file replication fails

CSM commands fail or CFM file replication fails. Refer to the applicable documentation that came with CSM and complete the problem resolution process.

rpower or rconsole commands fail

CSM `rpower` and `rconsole` commands fail. Refer to the applicable documentation that came with CSM and complete the problem resolution process.

GPFS checks

Performance problems

See the GPFS problem resolution and GPFS Performance White papers included with your software.

GPFS file system failure

See the GPFS problem resolution and GPFS Performance White papers included with your software.

SNMP monitoring

The service processor network, Ethernet switches, and Myrinet switch can be monitored using SNMP. All devices should be configured to send their SNMP traps to the management server. The management server should be configured to use `trapd` so that SNMP traps can be translated to a human readable form and added to the syslog.

Use the `lsnode -l` command to determine the host name for the Remote Supervisor Adapter card and the service processor name associated with the failing node. Use the telnet command, SSH command, or a Web browser to connect to the Remote Supervisor Adapter using the adapters host name, and select options to configure SNMP.

Configuring SNMP alerts from Myrinet

The Myrinet 2000 network in the Cluster 1350 is installed with adapter cards. One can use a graphical user interface, Mute, to monitor the entire network for events, which are logged and reported by the monitoring cards. You can configure an SNMP client or use a Web browser to access the monitoring card information. You can configure the monitoring cards to notify you of events by email.

The following Myrinet software packages are required:

- GM software is the base software required to use Myrinet 2000 network. It is the message-passing system for Myrinet networks, and includes a driver, Myrinet interface control program, a network mapping program, the GM API, library, and header files.
- m3-dist package, which provides the source for building the SNMP library for the GM layer.
- Mute (GUI) tool, which monitors the Myrinet network.

Use the following list to build the software:

- GM including the mt tools

- m3-dist (has dependency on GM)
- Mute (has dependency on GM and m3-dist)

Comprehensive details on how to build the software is described in the:

- Linux README file
- GM mt/README file
- m3-dist README file
- mute software README file

Currently m3-dist and Mute compile against GM version 1.5. With GM version 1.4 the SNMP library does not build m3-dist or Mute. Build the software using GM version 1.5.

You can access Myrinet software from: http://www.myri.com/scs/index.html (for GM, select the *LANai9* software).

Resetting the Remote Supervisor Adapter card

The Remote Supervisor Adapter (RSA) card is typically connected to a remote power control strip. To reset the RSA, plug the remote power control strip of the failing RSA card into another power source and issue the `power off` and `power on` commands to the RSA port.

Checking service processor logs

At the console prompt, type, `lsnode -l` to determine the host name for the Remote Supervisor Adapter (RSA) card and the service processor name associated with the node. At the console prompt, telnet, use SSH, or open a Web browser to connect to the RSA card using the host name, and select **View Log**.

Management, cluster, and storage node problems

The IBM components used for management, cluster, and storage nodes are shown in Table 20:

Table 20. IBM components used for management, cluster, and storage nodes

Node type	IBM component used
Management node	• xSeries 345 • @server 325 (64-bit operating system environment)
Cluster node	• @server 325 • xSeries 335 • xSeries 345 • BladeCenter
Storage node	• @server 325 • xSeries 335 (do not use in a Fibre-Channel storage configuration) • xSeries 345 • xSeries 360

While the @server 325, xSeries 335, BladeCenter, xSeries 345, and xSeries 360 are all high-reliability units, occasionally a component may fail. Two areas that might cause problems are:

1. Disk drives
2. System board

The following section includes information about:

- Disk drive failures
- System board failures
- xSeries 335 problems
- BladeCenter problems
- Power problems

Disk drive failures

The following section discusses issues with disk drive failures.

Disk drive failure on the management node

Use the following troubleshooting information about disk drive failures on the management node.

The xSeries 345 supports hot swapping of hard disks. To replace a failing hard disk on the management node:

1. Remove the failing hard disk by sliding it out through the front panel opening.
2. Slide the replacement disk drive into the open slot. The drive will rebuild automatically on a mirrored system. If the system is not mirrored a complete re-install of the management node is required. See Chapter 6, “Installing the software,” on page 31 for detailed instructions.

Disk drive failure on a cluster node

Use the following troubleshooting information about disk drive failures on the cluster node.

The xSeries 335 supports hot swapping of hard disks, but BladeCenter does not. To replace a failing hard disk on an xSeries 335:

1. Remove the failing hard disk by sliding it out through the front panel opening.
2. Slide the replacement disk drive into the open slot.
3. At the management node, issue the following command:

   ```
   installnode -n
   ```

   where *n* is the number of the node being reinstalled. If needed, have the customer contact support to assist with the correct naming conventions and IP addresses.

If a hard drive fails on a Blade server in the BladeCenter, first power down the Blade server. Next, remove the Blade server from the BladeCenter and replace the hard drive as outlined in *IBM eServer BladeCenter Hardware Maintenance Manual and Troubleshooting Guide*. Once the drive is replaced and the Blade server is returned to the BladeCenter issue the following command at the management node:

```
installnode -n
```

where *-n* is the number of the node being reinstalled. If needed, have the customer contact support to assist with the correct naming conventions and IP addresses.

Disk drive failure on a storage node

Use the following troubleshooting information about disk drive failures on a storage node.

The xSeries 345 and xSeries 360 support hot swapping of hard disks. To replace a failing hard disk on the storage node:

1. Remove the failing hard disk by sliding it out through the front panel opening.
2. Slide the replacement disk drive into the open slot. The drive will rebuild automatically on a mirrored system. To rebuild the storage node, if the system is not mirrored, enter the following command at the management node command line prompt:

   ```
   installstorage 1
   ```

System board failures

Use the following troubleshooting information about system board failures.

1. Replace the system board.
2. Flash the system BIOS to the level used in the installation. Refer to “Software version matrix” on page 31 for a listing of the software and firmware levels used in the Cluster 1350.
3. Flash the Diagnostics to match the BIOS level. Refer to “Software version matrix” on page 31 for a listing of the software and firmware levels used in the Cluster 1350.
4. Flash the onboard ASM to the current level. Refer to “Software version matrix” on page 31 for a listing of the software and firmware levels used in the Cluster 1350.
5. Perform the following configuration settings:
   - Devices and I/O Ports: `PORT 3F8, IRQ4`
   - Remote Console Redirection: Enabled, COM1, 9600, 8, None, 1, VT100, Enabled
   - Boot sequence: Diskette Drive, CD ROM, Network, Hard Drive 0, Boot Fail Count: DISABLED
   - Set the remote control password if a Remote Supervisor Adapter card is installed in this node (xSeries 335, xSeries 345, xSeries 360 only).
   - If you are replacing an HS20 BladeCenter with a serial port option, make sure that switch 7 in the switchblock is turned **on**.
   - Update the cluster software with the new MAC address associated with the new system board or Blade card you installed.
     - To get the MAC address of eth0 for the new component, use the CSM command: `getadapters -wn <nodename>`.

   Contact support for any setup or IP configurations that need to be performed.
6. Turn the customer over to support for any additional tasks needed to restore the node to full functionality.

xSeries 335 problems

In a xSeries 335 with an Remote Supervisor Adapter (RSA) over C2T connection make sure that the cluster node at the beginning of the C2T chain has an RSA card, external dongle, and connection to the onboard RSA processor.

BladeCenter problems

When the serial port option is used on a blade server in the BladeCenter it is important to make sure that switch number 7 in the switchblock is set to the ON position, that the card is fully seated in the option card slot, and that the cable is plugged into the serial header port. An improper switchblock setting, loose option card, or unplugged cable will cause the blade server to become unresponsive to `rconsole` commands.

CAUTION:
When two processors are installed, take special care not to pinch the cable under the metal standoff on the inside of the cover.

If the Ethernet Switch Module (ESM) is replaced in the BladeCenter then you must reassign the IP address for the external ports to work. Make sure the address is in the range reserved for the cluster LAN (.20 address) and not the management LAN.

Make sure that the PDUs in the cluster are connected to 220V source power. BladeCenters connected to PDUs plugged into 115V power will not function properly.

Power problems

The following section includes information about:

- No power to multiple devices
- No power to an individual device

No power to multiple devices

Use the following troubleshooting information about power failures to multiple devices.

1. Check that the 30 amp twist lock plugs are locked into the customer supplied receptacles.
2. Check the main power breakers at the customer breaker panel and make sure they are on.
3. Measure the voltage on the power out side of the Frame Power Block. If no voltage is present have the customer's electrician check for power issues. If no problems are found with the customer's power then replace the Input Power Block (FRU 32P1077). If the correct voltage is present, continue with the next step.
4. Make sure that the Power Distribution Unit (PDU) breakers are in the ON position.
5. Make sure that sure the PDU plugs are securely seated into the Power Out sockets on the Frame Power Blocks.
6. Make sure that voltage at the Power Out ports on the PDU using a Multimeter. If no power is present replace the PDU (FRU 9N9671). Otherwise, continue with the next step.
7. Swap out the power cable on the failing unit. If power LEDs do not light up on the failing unit, replace the power supply, or replace the complete unit if the power supply cannot be replaced.

No power to an individual device

Use the following troubleshooting information about power failures to an individual device.

1. Make sure that the PDU plugs are securely seated into the Power Out sockets on the Frame Power Blocks.
2. Make sure that the voltage at the Power Out ports on the PDU using a Multimeter. If no power is present replace the PDU (FRU 9N9671). Otherwise, continue with the next step.

3. Swap out the power cable on the failing unit. If power LEDs do not appear on the failing unit, replace the power supply or complete unit if the power supply cannot be replaced.

Related publications

Additional hardware maintenance and problem resolution information relating to the xSeries 335 and xSeries 345 was included with the Cluster 1350.

The documentation might be updated occasionally to include information about new features, or technical updates might be available to provide additional information that is not included with your cluster. These updates are available from the IBM Web site. Complete the following steps to check for updated documentation and technical updates:

1. Go to **http://www.ibm.com/pc/support/**.
2. In the **Learn** section, click **Online publications**.
3. On the “Online publications” page, in the **Brand** field, select **Servers**.
4. In the **Family** field, select **Clustering**.
5. Click **Continue** and select the online documents that best fit your needs.

Chapter 9. Configuring the KVM console switch

There are three possible KVM console-switch options for the Cluster 1350:

- IBM NetBAY 2x8 console switch
- IBM NetBAY Advanced Connectivity Technology Remote Console Manager (RCM)
- IBM NetBAY Local Console Manager (LCM)

Configuring the console switch after device replacement

1. Connect a laptop computer running a terminal emulation program (such as HyperTerminal) to the configuration port on the back panel of the console switch using a DB9 to RS232 null modem cable.
2. Configure the terminal settings to:
   - **9600 baud**
   - **8 bits**
   - **1 stop bit**
   - **no parity**
   - **no flow control**
3. Plug the supplied power cord into the back of the console switch and then into the power distribution unit (PDU) supplying power to the cabinet.
4. Turn on the power to the console switch. The power indicator on the front of the unit will blink for 30 seconds while the console switch performs a self-test. Approximately 10 seconds after it stops flashing, press the **Enter** key to access the main menu.
5. From the Terminal Applications menu, select **option 1** and set your network speed. Whenever possible, set your connection speed manually without relying on the auto-negotiation feature. After you have entered the selection, you will be returned to the *Network Configuration* menu.
6. Select **option 2** and specify if you are using a static or BootP IP address. Use a static IP address for ease of configuration. If you are using a BootP address, configure your BootP server to provide an IP address to the console switch and skip the next four steps.
7. From the Terminal Applications menu, select **option 3** and specify the IP address for the console switch.
8. From the Terminal Applications menu, select **option 4** and specify the Netmask for the console switch.
9. From the Terminal Applications menu, select **option 5** and type the default gateway address for the console switch.
10. To return to the main menu, enter **0**.

Use the following section to update the FLASH level on the console switch.

Upgrading the console switch FLASH level

To perform this update you will need a TFTP server. If you don’t already have a TFTP server, there are several you can download from the Internet. You must download the latest FLASH firmware from Avocent at http://www.avocent.com/support or copy the FLASH upgrade file (.fl file extension) from the CD shipped with the remote console switch. Save the FLASH upgrade file to an applicable directory on the TFTP server. After the FLASH upgrade file is copied, complete the following steps to upload the new FLASH file onto the console switch:

© Copyright IBM Corp. 2004

1. Connect a laptop computer running a terminal emulation program (such as HyperTerminal) to the configuration port on the back panel of the console switch using a RS232 DB9 null modem cable.
2. Set the console terminal to:
   - **9600 baud**
   - **8 bits**
   - **1 stop bit**
   - **no parity**
   - **no flow control**
3. Connect the LAN port in the console switch to an Ethernet hub being used as the TFTP server. Launch both the server software and the terminal emulation software.
4. Make sure that the power on the console switch is turned **on**. After approximately 40 seconds, the console switch will send out a message reading: `Avocent AutoView 1000R/2000R Ready_Press any key to continue`. Press any key to access the AutoView 1000R/2000R main menu.
5. Get the IP address of the TFTP server. If you are using the SolarWinds TFTP server, the IP address is shown in the lower right-hand corner of the server pane. Otherwise, you must extract the IP address using Windows operating system tools.
6. Right-click on **Network Neighborhood**, and select **Properties**.
7. Click the **Protocols** tab, and select **TCP/IP protocol**.
8. Select **Properties**, and make note of the IP address.
9. If necessary, assign the IP address for the console switch:
   a. To select the network configuration, in the terminal emulation window, type `1`.
   b. Compare the remote console switch IP address to the TFTP server IP address. The first 3 numbers of both IP addresses must be the same, but the last number must be unique. If the console switch IP address is incorrect, type `3` to select the IP address, and then enter the correct address.
   c. To exit network configuration, type `0` and follow the onscreen prompts to upgrade the FLASH level on the console switch.

Replacement of NetBAY Advanced Connectivity Technology Remote Console Manager

Detailed removal, replacement, and configuration information for the remote console manager (RCM) is addressed in the applicable service manual that came with the unit.

For more detailed information about configuration, removal, and replacement for the remote console manager, see http://www-306.ibm.com/pc/support/site.wss/

Replacement of NetBAY Advanced Connectivity Technology Local Console Manager

Detailed removal, replacement, and configuration information for the local console manager (LCM) is addressed in the applicable service manual that came with the unit.

For more detailed information about configuration, removal, and replacement for the local console manager, see http://www-306.ibm.com/pc/support/site.wss/

Chapter 10. Using the KVM switch

The KVM switch allows the use of a single keyboard, mouse, and monitor for multiple servers. You can switch between nodes and a console through the KVM switch interface, known as OSCAR.

The switch provides on-screen configuration and activity reporting, programmable scanning, NVRAM for saving configuration parameters, and an external reset switch.

Saving the KVM switch settings

Save device settings in the nonvolatile memory (NVRAM) of the KVM switch when any of the following occurs:

- The KVM switch is initially powered on.
- Nodes are added to or removed from the cabinet.
- There is a change in the keyboard, mouse, or monitor.

Attention: If device settings are not saved and the power to the KVM switch is lost, it may be necessary to reboot each node in the system to re-establish keyboard and mouse communications.

To save the device settings in the KVM switch NVRAM, perform the following steps:

1. On the keyboard, press **Print Screen**. The OSCAR selection window opens.
2. Press **F2**. The Advanced menu window opens. The Commands menu is highlighted.
3. Use the arrow keys (<- and ->) to highlight **Snapshot**, and press **Enter**. The device settings are now saved to NVRAM.

Connecting components with the KVM switch power turned on

You can connect additional servers to the KVM switch while the system is running. When you power up the newly connected node, the KVM switch recognizes it, and you can switch to the new node without taking any additional steps.

You can also connect the mouse, keyboard, and/or monitor to the KVM switch while the system is powered up. When you connect a new device, the KVM switch recognizes it and configures it to the settings of the currently selected node. This allows replacement of failed devices without having to restart the system.

Switching between nodes and the console

The KVM switch lets you disconnect the keyboard, mouse, and monitor from the currently selected node or from the console. You can also connect the keyboard, mouse, and monitor to another node or to the console.

Perform the following to switch between nodes or the console:

1. Press **Print Screen**. The OSCAR selection window opens.

   Attention: The servers and the console are listed in order by port or by name, depending on the user-definable settings in OSCAR menu attributes.
2. To select a node or the console, perform one of the following:
   a. Use the arrow keys (<- and ->) to select the node or the console; then press **Enter**.

© Copyright IBM Corp. 2004

   b. Press the numeric key that corresponds to the node port number or the console port number, then press **Enter**.
   c. Double-click the node or the console that you want to select.
3. Press the **Esc** key to exit OSCAR and close the OSCAR selection. The status flag window opens to indicate the currently connected node or the console.

Security features

The KVM switch provides for system security through the OSCAR interface. This security provides a simple keyboard and screen lock.

Open the security screen by selecting, **Advanced Menus > Setup > Security**. You must always provide a password to access the fields on the screen. The default password is **oscar**.

You can change passwords, set wait-time for locking to take effect, and set low-power mode for monitors so configured.

Resetting the mouse and keyboard

If the mouse and keyboard are not working properly (for example, no cursor response), you may need to reset the mouse and keyboard to restore the correct settings for the selected node. Perform the following steps to reset the mouse and keyboard:

1. Press the **Print Screen** key. The OSCAR selection window opens.
2. Press **F2**. The Advanced menu window opens. The Commands menu is selected.
3. Use the arrow keys (<- and ->) to select **Reset**, and press **Enter**. The mouse and keyboard are now reset.

If the mouse or keyboard are still locked up, you can push the reset button on the back panel to reset the KVM switch. Pressing the reset button might allow you to recover the device settings without unplugging and replugging the power cable for the node.

Chapter 11. Configuring a terminal server after device replacement

If you can successfully ping the iTouch IR-8000 series (20-port and 40-port), the In-Reach LX-4000 (32-port, 48-port) terminal server, no further action is needed. The terminal server has been properly configured.

Note: Do not define ports 21 through 40 on a terminal server with only 20 ports.

Configuring the ITouch-IR-8000 series and In-Reach LX-4000 series terminal server

If you cannot ping the terminal server, then configure the device:

1. Connect a laptop computer running a terminal emulation program (such as HyperTerminal) to the command port on the back panel of the terminal server using a DB9 to RJ45 serial cable.
2. Set the terminal configuration to the following settings:
   - **9600 baud**
   - **8 bits**
   - **1 stop bit**
   - **no parity**
   - **no flow control**
3. Attach a serial terminal to the command port. The default command port is the last port, either port 20 or port 40 depending on the size of the terminal server.
4. Turn on the terminal server and press **Enter** repeatedly until you see the `*Login>*` command prompt, and type `access`. No readable characters are visible.
5. Press **Enter**.
6. At the `*Username>*` command prompt, type `system` and press **Enter**. The In-Reach directory prompt displays.
7. At the `In-Reach>` command prompt, type `set priv` and press **Enter**.
8. At the `*Password>*` command prompt, type `system` .
9. At the `*In-Reach>*` command prompt, type `show ip` to see the current network settings.
10. To set the IP address, type `define ip address xxx.xxx.xxx.xxx`.
11. To set the gateway address, type `define ip primary gateway address xxx.xxx.xxx.xxx`.
12. To set the subnet mask, type `define ip subnet mask xxx.xxx.xxx.xxx`.
13. To save the configuration and restart the terminal server, type `init delay 0`.

If the terminal server does not already have an IP address, it might need further configuration so the serial ports can operate properly. Complete the following steps:

1. Telnet to the IP address assigned to the terminal server.
2. At the `*Login>*` command prompt, type `access`.
3. At the `*Username>*` command prompt, type `system`.
4. At the `*In-Reach>*` command prompt, type `set priv`.
5. At the `*Password>*` command prompt, type `system` .
6. At the `*In-Reach_Priv>*` command prompt, define the ports by entering the following:

```
In-Reach_Priv>define port 1-20 access remote
In-Reach_Priv>define port 21-40 access remote
In-Reach_Priv>define port 1-20 flow control enable
In-Reach_Priv>define port 21-40 flow control enable
```

© Copyright IBM Corp. 2004

```
In-Reach_Priv>define port 1-20 speed 9600
In-Reach_Priv>define port 21-40 speed 9600
In-Reach_Priv>define port 1-20 que disable
In-Reach_Priv>define port 21-40 que disable
In-Reach_Priv>lo port 1-20
In-Reach_Priv>lo port 21-40
In-Reach_Priv>init delay 0
```

The last command saves any configuration changes and restarts the terminal server. The terminal server should now be fully operational.

Related publications

For more information about the In-Reach LX-4000 terminal servers configuration settings, see http://service.mrv.com/support/index.cfm.

Chapter 12. Configuring and replacing the Cisco 3550 10/100 Ethernet switch

To replace the Cisco 3550 10/100 (24-port or 48-port) Ethernet switch, see the instructions that came with your switch.

Configuring after device replacement

To set up the new 10/100 Ethernet switch you will need the following:
- Laptop computer
- Hyper Terminal program
- DB9 to RJ45 serial cable

Take the following steps:

1. Connect the RJ45 side of the serial cable to the port on the front of the Cisco switch marked CONSOLE.
2. Connect the other end of the cable to the laptop computer.
3. Start the Hyper Terminal application. Configure the terminal to:
   - **9600 baud**
   - **8 bits**
   - **no parity**
   - **1 stop bit**
   - **no flow control**
   - **VT100 Emulation**
4. At the command prompt in the terminal emulation window, type: **enable**. This command enables the administrative mode.
5. At the command prompt, type **ibm** and press **Enter**. The prompt will change from a **>** to a **#** to indicate you are in administrative mode.
6. At the command prompt, type **show run** to show the current configuration information. Make note of the current settings.

   Collect the following information to set up the new switch:
   - Switch IP address
   - IP mask
   - Default gateway IP address
   - Switch host name
   - Cluster name
7. At the **#** prompt, enter the following commands:

   ```
   configure terminal
   interface vlan1
   ip address 172.xxx.xxx.xxx
   255.255.xxx.xxx
   exit
   ip default-gateway 172.xxx.xxx.xxx
   end
   show run-config
   copy running-config startup-config
   ```
8. Type **exit** to log out of the terminal session.

© Copyright IBM Corp. 2004

The setup procedures are documented in the Cisco Quick Start Guide Catalyst 3500 Series XL Switches: http://www.cisco.com/univercd/cc/td/doc/product/lan/c2900xl/29_35xu/3500.htm

After completing the First Time Setup section in the Cisco Quick Start guide, save it to the *startup* file so the switch can be rebooted without losing the setup configuration. At the `telnet` prompt, type the command:**`copy run start`**

The Quick Start guide also describes how to obtain the JAVA plug-in and configure your browser to support the HTML interface.

There is an SNMP vulnerability for various versions of switch firmware.

See http://www.cisco.com/warp/public/707/cisco-malformed-snmp-msgs-pub.shtml for specific firmware patches to download.

Setup troubleshooting

Once the initial setup is complete, there should be a network connection between the PC and the switch. If a ping to the switch fails, make sure that the IP address and gateway address to make sure the subnet and gateway addresses match:

- On the PC, at the command prompt, type: `ipconfig`
- On the switch, at the command prompt, type: `show running x`

Related publications

Catalyst 5000 Family Ethernet and Fast Ethernet Switching Modules Installation and Configuration Note (including Translated Safety warnings 10 languages): http://www.cisco.com/univercd/cc/td/doc/product/lan/cat5000/cnfg_nts/ethernet/5014etsm.htm#20508

Catalyst 3500 Series XL Hardware Installation Guide Includes Troubleshooting: http://www.cisco.com/univercd/cc/td/doc/product/lan/c2900xl/29_35wc/3500ig/index.htm

Catalyst 2900 Series XL and 3500 Series XL Cisco IOS Release 12.0(5.3)XU: http://www.cisco.com/univercd/cc/td/doc/product/lan/c2900xl/29_35wc/rn53/1061505.htm

Quick Start Guide Catalyst 3500 Series XL Switches: http://www.cisco.com/univercd/cc/td/doc/product/lan/c2900xl/29_35xu/3500.htm

Chapter 13. Configuring and replacing a Cisco Catalyst 3750 Gigabit Ethernet switch

Before you install, operate, or service the system, read the Cisco *Site Preparation and Safety Guide*. This document contains important safety information you should know before working with the system.

Replacing the Cisco Catalyst 3750 Gigabit Ethernet switch

Detailed hardware maintenance information covering installation, removal, and replacement procedures for the Cisco 3750 24-port stackable Gigabit Ethernet switch are found at http://www.cisco.com/univercd/cc/td/doc/product/lan/c2900xl/29_35wc/3500ig/index.htm.

Configuring after device replacement

To set up the new switch you need the following:

- Laptop computer
- Hyper Terminal program
- DB9 to RJ45 serial cable

Take the following steps:

1. Connect the RJ45 side of the serial cable to the port on the front of the Cisco switch marked **CONSOLE**.
2. Connect the other end of the cable to the laptop computer.
3. Start a console terminal, such as HyperTerminal. Configure the terminal settings to:
   - **9600 baud**
   - **8 bits**
   - **no parity**
   - **1 stop bit**
   - **no flow control**
   - **VT100 emulation**
4. At the command prompt in the terminal emulation window, type `enable`. This will put you in administrative mode.
5. At the prompt, type `ibm` and press **B**. The prompt will change from a **>** to a **#** to indicate you are in administrative mode.
6. Type `show run` to show the current configuration information. Make note of the current settings and then logoff from the session.

You need the following information to set up the new switch:

- Switch IP address
- IP mask
- Default gateway IP address
- Switch host name
- Cluster name

© Copyright IBM Corp. 2004

Use the following information to set up the switch after its been replaced.

Cisco Catalyst 3750 Gigabit Ethernet switch - The set up procedures for the Catalyst 3750 24-port stackable Gigabit Ethernet switch are at http://cisco.com/en/US/products/hw/switches/ps5023/prod_technical_documentation.html.

Setup troubleshooting

Once the initial setup is complete, there should be a network connection between the PC and the switch. If a `ping` to the switch fails, make sure that the IP address and gateway to make sure the subnet and gateway addresses match:

- On the PC, at the command prompt, type: **`ipconfig`**
- On the switch, at the command prompt, type: **`show running`**

Nodes on the same VLAN can communicate via `ping and telnet`. They cannot communicate to nodes on different VLANs. Make sure that VLANs:

- Connect node1 and node2 to the same VLAN and `ping` node2 from node1. This ping should succeed.
- Connect node1 to VLAN1 and node2 to VLAN2 and `ping` node2 from node1. This ping should fail.

Related publications

Quick Start Guide Catalyst 3500 Series XL Switches: http://www.cisco.com/univercd/cc/td/doc/product/lan/c2900xl/29_35xu/3500.htm

Cisco Catalyst 3750 Switch Hardware Guide at http://www.cisco.com/univercd/cc/td/doc/product/lan/cat3750/12114ea1/3750hig/index.htm

Chapter 14. Configuring and replacing the Cisco 4000 series switch

Installation, removal, replacement, and troubleshooting procedures

Detailed hardware maintenance information covering installation, removal, and replacement procedures for the Cisco 4000 series switch is found at http://www.cisco.com/univercd/cc/td/doc/product/lan/cat4000/hw_doc/install/

Detailed troubleshooting procedures for the Cisco 4000 Series switch are found at http://www.cisco.com/univercd/cc/td/doc/product/lan/cat5000/trbl_ja.htm

Additionally, IBM has included with each Cisco 4000 Series switch a specially designed thermal duct to make sure proper airflow around the switch. Figure 5 shows an exploded view of the thermal duct and how it fits within the cabinet and attaches to the switch.

Figure 5. Thermal duct used with Cisco 4000 switches

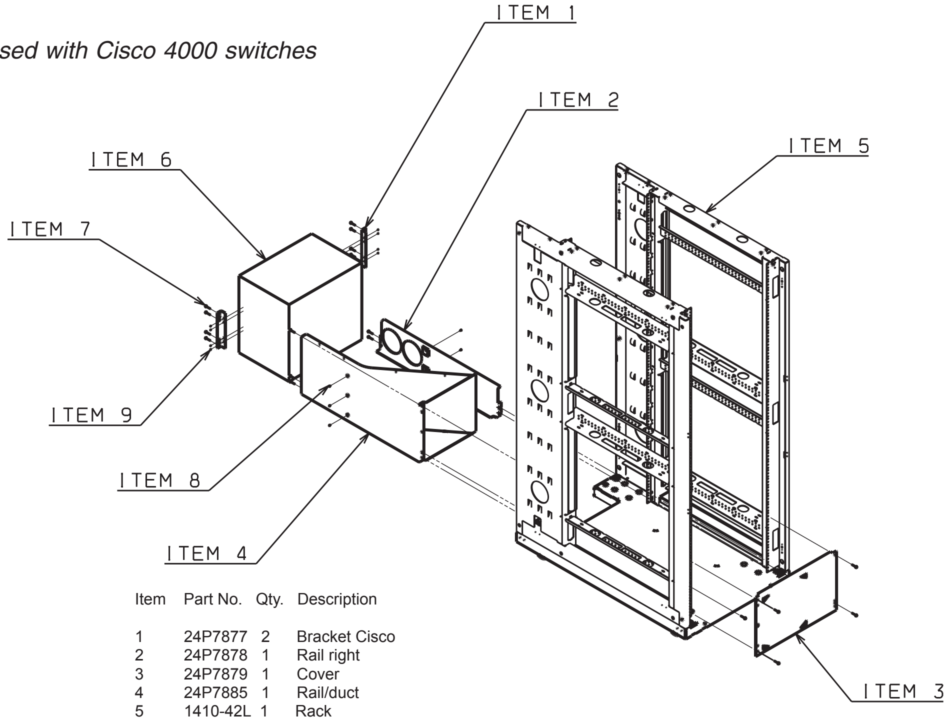

Item	Part No.	Qty.	Description
1	24P7877	2	Bracket Cisco
2	24P7878	1	Rail right
3	24P7879	1	Cover
4	24P7885	1	Rail/duct
5	1410-42L	1	Rack
6	Ref. only	1	Cisco 400x switch
7	12J5289	1	M6 Hex flange screw
8	1621811	5	M4 Hex flange screw
9	N/A	6	M4 Flat head screw

If the Cisco 4000 Series switch is ever removed for maintenance make sure the thermal duct is reinstalled whenever the switch is returned to the cabinet. Failure to reinstall the thermal duct could create temperature management problems within the cabinet.

Related publications

Additional information on a variety of topics (including software configuration) for the Cisco 4000 series switches is available at http://www.cisco.com/univercd/cc/td/doc/product/lan/cat4000/

© Copyright IBM Corp. 2004

Chapter 15. Configuring and replacing Cisco Catalyst 6500 series devices

Before you install, operate, or service the system, read the Cisco *Site Preparation and Safety Guide*. This document contains important safety information you should know before working with the system.

Additionally, IBM has included with each Cisco 6503 10/100/1000 Mb Ethernet switch rack cabinet and each Cisco Catalyst 6509 10/100/1000 Mb Ethernet switch rack cabinet a specially designed rack-mounting kit with thermal air duct to make sure that there is proper airflow around the switch.

Replacing the Cisco Catalyst 6503 and 6509 switch rack unit and rack-mount assemblies

Use the following figures to help you replace a Cisco Catalyst 6503 or 6509 10/100/1000 Mb Ethernet switch rack unit.

Detailed hardware maintenance information covering installation, removal, and replacement procedures for the Cisco 6503 and Cisco 6509 switch rack units are found at http://www.cisco.com/uivercd/cc/td/doc/product/lan/cat6000/6000hw/index.htm.

Figure 6 shows an exploded view of the Cisco Catalyst 6503 switch rack unit and rack-mounting rail assembly and how it fits within the cabinet and attaches to the switch.

Figure 6. Cisco Catalyst 6503 rack mounting kit with air duct

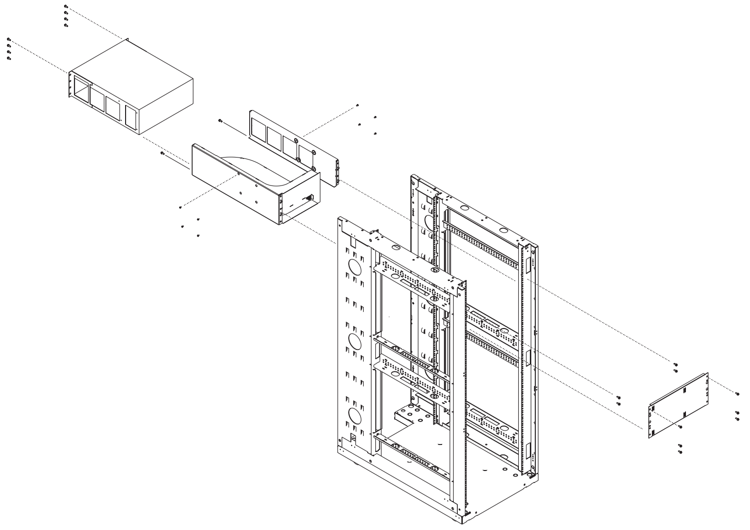

© Copyright IBM Corp. 2004

Figure 7 shows an exploded view of the Cisco Catalyst 6509 switch rack unit and rack-mounting rail assembly and how it fits within the cabinet and attaches to the switch.

Figure 7. Cisco Catalyst 6509 rack-mounting kit with air duct

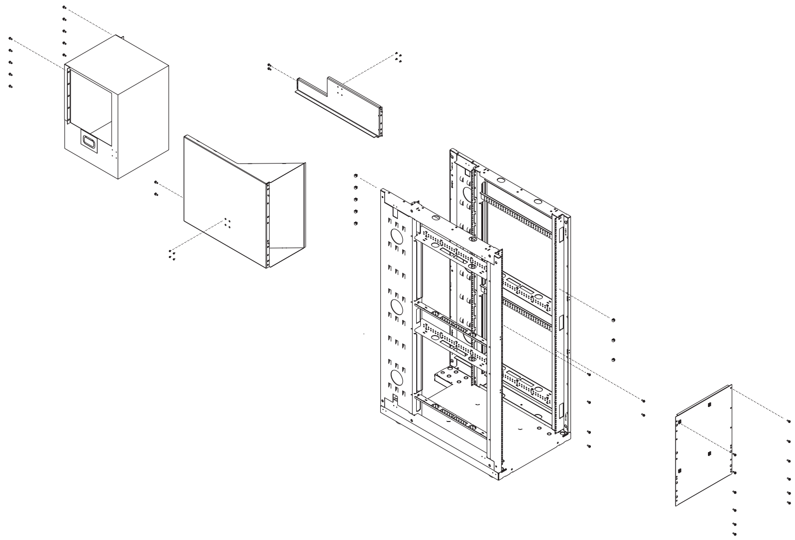

Replacing the Cisco Catalyst Ethernet line cards and 10 Gigabit switch

Detailed hardware maintenance information covering installation, removal, and replacement procedures for the Cisco 6548 and Cisco 6748 Ethernet line cards, and the Cisco 6704 Gigabit Ethernet switch are found at http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/index.htm.

Configuring after device replacement

To set up the new switch you need the following:

- Laptop computer
- Hyper Terminal program
- DB9 to RJ45 serial cable

Take the following steps:

1. Connect the RJ45 side of the serial cable to the port on the front of the Cisco switch marked **CONSOLE**.
2. Connect the other end of the cable to the laptop computer.
3. Start a console terminal, such as HyperTerminal. Configure the terminal settings to:
   - **9600 baud**
   - **8 bits**
   - **no parity**
   - **1 stop bit**
   - **no flow control**
   - **VT100 emulation**
4. At the command prompt in the terminal emulation window, type `enable`. This will put you in administrative mode.

5. At the prompt, type **ibm** and press **B**. The prompt will change from a **>** to a **#** to indicate you are in administrative mode.
6. Type **show run** to show the current configuration information. Make note of the current settings and then logoff from the session.

You need the following information to set up the new switch:
- Switch IP address
- IP mask
- Default gateway IP address
- Switch host name
- Cluster name

Setup troubleshooting

Once the initial setup is complete, there should be a network connection between the PC and the switch. If a `ping` to the switch fails, make sure that the IP address and gateway to make sure the subnet and gateway addresses match:
- On the PC, at the command prompt, type: **ipconfig**
- On the switch, at the command prompt, type: **show running**

Nodes on the same VLAN can communicate via `ping and telnet`. They cannot communicate to nodes on different VLANs. Make sure that VLANs:
- Connect node1 and node2 to the same VLAN and `ping` node2 from node1. This ping should succeed.
- Connect node1 to VLAN1 and node2 to VLAN2 and `ping` node2 from node1. This ping should fail.

Related publications

Cisco Catalyst 6500 series switch documentation: http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/index.htm.

Chapter 16. Configuring and replacing the SMC 8624T 10/100/1000 switch

To replace the SMC 8624T (24-port) 10/100/1000 switch, see the instructions that came with your switch. See http://www.smc.com/index.cfm?action=work_home.

Configuring after device replacement

To set up the new switch you will need the following:

- Laptop computer
- Hyper Terminal program
- DB9 to DB9 (DTE pinout) serial cable

Take the following steps:

1. Using a DB9 serial cable, connect a serial console, running a terminal emulation program (such as, Hyperterminal), to the serial port on the switch.
2. Define the configuration settings to the following:
   - **9600 baud**
   - **8 bits**
   - **no parity**
   - **1 stop bit**
   - **no flow control**
   - **VT100 emulation**
3. At the command prompt, in the terminal emulation window, type: Type `admin` as the user name and press **Enter**.
4. Type `admin` as the password and press **Enter**. This puts you in privilege executive mode.

   Collect the following information to set up the new switch:
   - Switch IP address
   - IP mask
   - Default gateway IP address
   - Switch host name
   - Cluster name
5. At the command prompt, type `copy running-config startup-config` and press **Enter**. This copies the changes to the startup-config file so that all modifications are saved when the switch is restarted.
6. Type `exit` to exit the privilege executive mode.
7. Type `exit` to close the terminal session.

Related publications

Detailed hardware maintenance information covering installation, removal, troubleshooting, and replacement procedures for the SMC 8624T Gigabit switch is found in the *SMC 8624T Management Guide* at http://www.smc.com/drivers_downloads/library/smc8624t1b.pdf

Detailed firmware update instructions, diagnostics, utility software, and user guides are found at http://www.smc.com/index.cfm?action=products_downloads&productCode=SMC

© Copyright IBM Corp. 2004

Chapter 17. Configuring and replacing the Myrinet 2 Gigabit switch

The 2 Gb Myrinet switch is an option that provides high-speed communication between the storage nodes and cluster nodes, over an optical cable. It requires a Myrinet switch chassis in the primary cabinet and a Myrinet PCI adapter in each storage node and cluster node.

Myrinet PCI adapter

The Myrinet PCI adapter resides in the cluster and storage nodes. Use the installation procedures in the applicable server documentation to replace a Myrinet PCI adapter.

The GM software for running the Myrinet adapter resides in the cluster and storage nodes, so no new installation of software is required when a Myrinet PCI adapter is replaced.

Myrinet switch cabinet

The Myrinet switch cabinet contains the following replaceable components:

8-port line card
Provides the connections to the storage and cluster nodes. The line cards plug into slots in the switch chassis.

Management Module
Manages and routes the Myrinet traffic, polling the ports and building tables to control the addressing of messages.

Blower module
Cools the Myrinet switch chassis

All of these components can be hot-swapped. The Myrinet documentation discusses installation of these components.

The three Myrinet chassis sizes available are described in Table 21.

Table 21. Myrinet chassis sizes and capacities

Slots in switch	Line cards	Nodes supported	EIA slots consumed
5	1-4	4-32	4
9	1-8	4-64	6
17	1-16	4-128	10

If the backplane fails in the Myrinet switch, you must replace the entire switch chassis. Use the following steps to replace the chassis:

1. Make sure that the cluster is not running critical applications.
2. If the optical cables connected to the switch are not labeled, place labels on the cables so they can be located to their respective connectors when the new chassis is installed.
3. Disconnect the optical cables from the connectors on the Myrinet switch. You do not need to power down or change the configuration of the switch before doing this.

© Copyright IBM Corp. 2004

Note: Be sure to install dust caps on all the connectors after the cables are removed.

4. Disconnect the power cord from the Myrinet switch. This powers down the switch.
5. Remove the rack-mount screws from the chassis; then remove the chassis from the rack.
6. Install the new chassis and fasten the rack-mount screws.
7. Connect the optical cables to the connectors on the switch.

   Note: Save the dust caps for future use.
8. Connect the power cord to the Myrinet switch. This powers up the switch.

Configuring the switch after device replacement

The Myrinet switch automatically remaps all the PCI boards, so no manual configuration is needed.

IBM Customer Support personnel will update the firmware if necessary.

Related publications

Additional installation and troubleshooting information is available online from Myricom at the following URL: http://www.myri.com/scs/#documentation

Chapter 18. Configuring and replacing a Topspin device

The Cluster 1350 supports the:

- Topspin 120 InfiniBand switch - provides either 24-ports of 10 Gbps connectivity or 8-ports of 30 Gigabits per second (Gbps) connectivity in a single 1-U unit.
- Topspin InfiniBand host channel adapter - provides two 10 Gbps ports in a PCI short form-factor card.

Before you install, operate, or service the system, read “Safety” on page ix. This chapter contains important safety information you should know before working with the system.

Replacing the Topspin 120 InfiniBand switch

To replace the InfiniBand switch, complete the following steps:

1. Turn off the power to the cluster.
2. Locate the switch slot and remove the screws attaching it to the rack unit.
3. Slide the switch out of the rack unit.
4. Remove the InfiniBand cable connector:
   - pinch connector - Pinch both sides of the back of the connector and pull the connector away from the port.
   - pull connector - grasp the connector with one hand and push on the port; then, pull the latch away from the port with your other hand and gently wiggle the connector away from the port.
5. Unpack the switch, following static precautions on page xv, and slide it into the empty slot.
6. Reconnect the cable and replace and tighten the screws.

Replacing the Topspin Infiniband host channel adapter

To replace the host channel adapter, complete the following steps:

1. Turn off the power to the cluster.
2. Locate the host channel adapter and remove any screw attaching it to the rack unit.
3. Slide the adapter out of the slot.
4. Remove the InfiniBand cable connector:
   - pinch connector - Pinch both sides of the back of the connector and pull the connector away from the port.
   - pull connector - grasp the connector with one hand and push on the port; then, pull the latch away from the port with your other hand and gently wiggle the connector away from the port.
5. Unpack the replacement adapter, following static precautions on page xv, and slide it into the empty slot.
6. Reconnect the cable and replace and tighten the screws.

© Copyright IBM Corp. 2004

Configuring the Topspin 120 InfiniBand switch and host channel adapter after replacement

If you can successfully ping the InfiniBand switch or host channel adapter, no further action is needed. The devices are properly configured.

To set up the new InfiniBand switch and/or host channel adapter you will need the following:
- Laptop computer
- Hyper Terminal program
- DB9 to RJ45 serial cable

To configure the InfiniBand switch and host adapter, complete the following steps:

1. Using an RJ45 serial cable, connect a serial console, running a terminal emulation program (such as, Hyperterminal), to the serial port on the InfiniBand switch.
2. Define the configuration settings to the following:
   - **9600 baud**
   - **8 bits**
   - **no parity**
   - **1 stop bit**
   - **no flow control**
3. Attach the power cables, power on the InfiniBand switch. The InfiniBand switch automatically starts up.
4. After the start up sequence ends, press **Enter** repeatedly until the console displays the `Username:` command prompt.
5. Type `super` as the user name and press **Enter**.
6. At the `Password:` command prompt, type `super` and press **Enter**.
7. At the `Topspin-120` prompt, type `enable` and press **Enter**.
8. Type `configure` and press **Enter**.
9. To enter the config-if-mgmt-ethernet mode, type `interface mgmt-ethernet` and press **Enter**.
10. Type the IP address of the management port and netmask and press **Enter**.
11. Type the default gateway IP address and press **Enter**.
12. To enable the management port, type `no shutdown` and press **Enter**.
13. Save the configuration and type `exit` to exit the config-if-mgmt-ethernet mode and press **Enter**.
14. Type `exit` to exit the config mode and press **Enter**.

Installing device drivers and kernel software

An integrated suite of device drivers is provided with your InfiniBand switch and InfiniBand host channel adapter.

To detect an available kernel and install the device drivers:

1. Insert the Topspin 120 CD into the CD-ROM drive.
2. At the prompt, type `copy running-config startup-config mount /mnt/cdrom ./tsinstall` and press **Enter**.

3. After the installation completes, type **`ifconfig ib`**, followed by a corresponding IP address and **`netmask`**, followed by a corresponding IP address and press **Enter**. For example: `# ifconfig ib0 xxx.xxx.x.x netmask xxx.xxx.xxx.x`

For more information on configuring device drivers, see the *Topspin 120 InfiniBand switch Installation Guide* and the *Topspin 120 InfiniBand Host Channel Adapter Installation Guide* provided with the device.

Related publications

Detailed hardware maintenance information covering installation, removal, troubleshooting procedures, and replacement procedures for the Topspin 120 InfiniBand switch and Topspin Infiniband host channel adapter are found at http://www.topspin.com/solutions/products.html.

Chapter 19. Removing and replacing the Power Distribution Unit

The Power Distribution Unit (PDU) provides AC power within the cabinet. The PDUs are mounted sideways beside the regular rack space. Two types of PDUs are used:

- Rack PDUs
- Front-end PDUs

Rack PDUs provide power to components within a cabinet, while front-end PDUs provide the connection to the external power source and distribute the power among the rack PDUs. A rack PDU can also be directly connected to the external power source to eliminate the need for the front-end PDU. Up to four front-end PDUs can be placed in each cabinet and up to twelve rack PDUs.

To remove the Power Distribution Units, perform the following steps:

1. Shut down all devices.
2. Remove the side cover on the side of the rack that the failing PDU is located on.
3. Turn off each rack PDU using the breaker switch.
4. Unplug each rack PDU from the front-end PDU or customer supplied power source.
5. If present, unplug the front-end PDU from the customer supplied power source.
6. Remove the four screws holding the plate on which the PDUs are mounted.
7. Turn the plate over to access the screws that hold the rack PDUs and the front-end PDU (if present) on the plate.
8. Remove the screws holding the failing component to the plate.
9. Replace the failing component (front-end PDU or rack PDU) and reverse the steps shown above to reinstall the PDUs.

© Copyright IBM Corp. 2004

Appendix A. Getting help and technical assistance

This appendix contains information about where to go for additional information about IBM @server Cluster 1350 and what to do if you experience a problem with it.

Before you call

Before you call, make sure that you have taken these steps to try to solve the problem yourself:

- Use the troubleshooting information in this document to try to resolve the problem.
- Check for updated technical information, hints, tips, or new device drivers at the @server Cluster 1350 InfoCenter Web site. Go to http://publib.boulder.ibm.com/cluster/current.htm.
- Check common answers to questions about IBM clusters. To subscribe to the CSM mailing list, go to http://www-124.ibm.com/developerworks/oss/mailman/listinfo/csm. To subscribe to the XCAT mailing list, go to http://www.xcat.org.
- If you suspect a software problem, see the information for the operating system or program.
- If you still experience a problem, contact Hardware Service and Support (see below). Be sure to have the following information available when you call.

Machine type: **1410** (or 1417-11LX)
Model: **42L** (42X or 25X)
Serial number:

- The label containing the serial number can be found on the purchase order or in the rack cabinet.
  - The 1410 models 42L and 42X rack model and serial number label is located in the rear, on the left rack panel, near the 26 U slot. You can have cabling obstructing the view of the label. Carefully move the cables apart to read the serial number.
  - The 1410 model 25X rack model and serial number label is located in the front, on the bottom-right of the rack unit, near the 1 U slot.
  - The 1417 model 11X rack model and serial number is generally located in the rear, on the bottom-right**.**

Hardware service and support

You can receive hardware service through IBM Integrated Technology Services or through your IBM reseller, if your reseller is authorized by IBM to provide warranty service. Go to http://www.ibm.com/planetwide/ for support telephone numbers, or in the U.S. and Canada, call 1-800-IBM-SERV (1-800-426-7378).

In the U.S. and Canada, hardware service and support is available 24 hours a day, 7 days a week. In the U.K., these services are available Monday through Friday, from 9 a.m. to 6 p.m.

© Copyright IBM Corp. 2004

Appendix B. Frequently asked questions

Here are some frequently asked questions about the IBM @server Cluster 1350.

Q: Why do I sometimes get the error message *"2651-689 Java interface error for method "query": SPException"?*

A: This is due to a defect in the Remote Supervisor Adapter (RSA) firmware that is currently being investigated by xSeries development. This problem occurs after making between 100 and 200 connections to the RSA through the ASM library. The work around is to reset the RSA using the web or telnet interface. Until this defect is fixed, you may want to increase the polling interval for each hardware control point using the command: `chrsrc -s 'Name like "%"' IBM.HwCtrlPoint PollingInterval=86400`

Q: When I issue an `'rpower -n <node> reboot` command why does the node not reboot?

A: Sometimes the Remote Supervisor Adapter (RSA) cards get hung. They can be reset via the web interface, telnetting to the RSA card, or issuing the command `rpower -n <node> resetsp_hcp`.

Q: During installation process tftp hangs on the installing node. What's going on?

A: tftp is not loaded/configured on the management node.

Q: Why doesn't the xSeries 345 boot PXE correctly?

A: You cannot have a PCI ethernet card that uses the e1000 driver in the xSeries 345 when installing. Take the card out and retry the installation.

Q: Why does the dhcp server run out of leases?

A: The problem may be that you have two networks going to the same switch fabric. This causes both *eth0* and *eth1* to see the dhcp requests. To fix this create separate VLANs in the switches, one for each network attached to the switch.

Q: Why did PXE boot not get an IP address using dhcp, but the operating system can?

A: Check the switch. All ports connected to nodes (management, compute, and storage) should have spinning-tree turned off.

Q: Why doesn't the storage node see the drives on the FastT700, but the orange light on the host adapter card still blinks?

A: The qla2300 driver did not load properly. Make sure the proper version of the driver is installed.

Q: What is causing SLES to continuously install the nodes?

A: Check that fully qualified names (host.domainname) are used in the */etc/hosts* file and that the command `dnsdomainname` returns the correct domain name. Also make sure that */etc/dhcpd.conf* file contains the line: *'option domain-name "cluster.net";'* Once these changes have been made run the `csmsetupyast`command

© Copyright IBM Corp. 2004

and then rerun the `installnode` command. If the install still cycles then edit the */csminstall/Linux/SIS/scripts/<hostname>.sh* file and comment out the shutdown line near the bottom of the file. Now using the console watch the boot and the error messages should be on the console when the process has completed.

Q: Why do SLES installs take forever?

A: Issue the `installnode` command and then on the management node immediately edit the */tftpboot/pxelinux.cfg/AC** files. Take out *`console= portion`* from the APPEND line. Now all messages will go to the KVM console and the install will be quicker.

Q: Why do SLES installs fail but issue no error message?

A: Modify the */tftpboot/pxelinux.cfg/<HEX>*

Where <HEX> is the IP address in hex format. If syslinux 2.00 or later is installed, you can use the `gethostip -x` command to get the HEX name of a node.

Appendix C. Error and event logs

There are multiple log files available to help monitor and troubleshoot the cluster:

Linux log

The Linux OS log can be viewed in */var/log/messages*

The system logging daemons are *syslogd* and *klogd*. They are configured via */etc/syslog.conf*.

Log files are automatically rotated by the `logrotate` command. To rotation is configured with the */etc/logrotate.conf* file.

Node log

PC Doctor 2.0 is a ROM-based Diagnostic resident on the servers made available by selecting F2 on boot up. PC Doctor error logs are in the diagnostic portion of the boot up. Press F2 to run diagnostics, then F3 to view log file.

POST/BIOS errors can be read by pressing F1 key during boot process and then selecting View Error Logs from menu. This gives a POST code and description of the error. For example:

```
301 Keyboard Input Error 164 Memory size has changed
```

Cluster System Management log

Cluster System Management (CSM) log files can be viewed in the */var/log/csm/installnode.log* file using the `reventlog -n <nodename> all` command.

Remote Supervisor Adapter log

Remote Supervisor Adapter (RSA) Adapter log files can be viewed by using `telnet` into the adapter and selecting the *View Log File* from the menu.

American Power Conversion log

You can view the American Power Conversion (APC) event log via Web, FTP or local console I/F:

1. `Telnet` to the switch.
2. From main menu, you will see CTL-L for Event Log.
3. Events are logged in descending order by date, time and event.

Linux Cluster Installation Tool event log

You can view the Linux Cluster Installation Tool (LCIT) event log through a local or remote console or terminal window:

1. Open a console window on the management node.
2. From the root directory, type: `tail -f /var/log/messages.`

© Copyright IBM Corp. 2004

Appendix D. Known problems

This chapter describes known Cluster 1350 problems and resolutions. Use this chapter to troubleshoot hardware errors.

Node

Amber light on node

There is an amber warning light on the node to indicate the log file is either at 75% or 100% full. To turn off the LED, clear the log.

There is a setting to wrap the log file so the LED never registers if the file is full:
1. Boot using the xSeries 335 Service Processor Firmware diskette.
2. In the Main Menu, select **Configuration Settings**.
3. In the Configuration Menu, select **General Settings**.
4. Set the `75% Full and Log Full` setting to **No**.

COM port settings in BIOS

The COM Port settings for the cluster node should be:

COM Port 1/A
2E8

COM Port 2/B
2F8

Move the serial port jumper from port A to port B on cluster nodes.

CSM

Stale NFS mounts

Existing NFS mounted file systems are inaccessible after a CSM installation on a cluster node.
1. Remount the NFS file systems.
2. If there is an existing */tftpboot* partition on cluster nodes, an error is displayed on the console during CSM installation on the cluster node. Even though an error is displayed, the CSM installation was still successful

rpower hard shut down

The `rpower` command performs a hard shut down. To shut down the OS prior to issuing the `rpower` command issue the following command:

```
dsh -a '/sbin/init 0'
```

Storage

Driver module ordering

During a standard install on the storage nodes the system will attempt to boot from disk located in the FAStT storage device connected to the Qlogic Fibre Channel (FC) Controller instead of the local SCSI drive connected to the internal Adaptec SCSI controller. Why this happens is as follows:

© Copyright IBM Corp. 2004

When the modules are loaded, the order ends up in such a way that the driver for the Qlogic FC controller gets loaded before the driver for the Adaptec SCSI Controller. This causes the probing for the devices to occur such that the Fabric gets assigned *sda*, *sdb*, and so on followed by the local SCSI disks. Make the following modifications to make sure that the SCSI module is loaded before the Fibre module. This will validate that the probing and naming assigns the *sda* device to the first local disk.

1. First, modify the */etc/modules.conf* file by adding the line, `options scsi_mod max_scsi_luns=128` to the end of *modules.conf*. Also remove unnecessary information and reorder the way the modules are loaded. An example of an edited file is as follows:

   Original **modules.conf**:

   ```
   alias eth0 e1000
   alias scsi_hostadapter qla2x00
   alias scsi_hostadapter1 aic7xxx
   alias scsi_hostadapter2 ips
   alias parport_lowlevel parport_pc
   alias scsi_hostadapter2 qla2x00
   alias usb-controller usb-ohci
   alias scsi_hostadapter4 aic7xxx
   ```

   Edited **modules.conf**:

   ```
   alias eth0 e1000
   alias scsi_hostadapter aic7xxx
   alias scsi_hostadapter1 ips
   alias eth1 e1000     alias eth1 e1000
   alias parport_lowlevel parport_pc
   alias scsi_hostadapter3 aic7xxx
   options scsi_mod max_scsi_luns=128
   ```
2. Next, rebuild the two **initrd** images:

   mkinitrd `initrd-2.4.2-2.img 2.4.2-2 -f`

   mkinitrd `initrd-2.4.2-2smp.img 2.4.2-2smp -f`
3. Reboot the node.

KVM

GUI does not appear on first node

If the GUI display does not appear on the first node of the C2T chain, use the text mode.

2x8 Switch powers on with console port B

To remedy this go into the menu settings and change from cooperative to preemptive mode, reselect port 2 and console A will appear. When working properly do a Snapshot to save the setting.

Cluster port 1 reboots

The cluster Port 1 may reboot on power up, and either boots up in text mode blinking every 5 seconds or boots up with a white screen. There are two methods to remedy this situation:

1. Manually select the other ports in the C2T string, then reselect node 1.
2. Unplug the server connections from the port, reattach them in order, and re-plug in the server.

Subsequent KVMs unresponsive

Make sure the KVM switch that was added is in default settings mode.

RSA and Service Processor

If there are any Remote Supervisor Adapter (RSA) errors, check to make sure the RSA is in PCI slot 2.

RSA unable to load firmware

This condition is indicated by error FFFF, 0007. Power cycle the RSA adaptor to clear this condition. The RSA may need to be replaced if this condition persists.

RSA/Service Processor invalid naming

There cannot be any spaces when assigning names of the RSA and Service Processor. If a name is not recognized, make sure that there are no trailing blanks after the name.

Light path points to PCI LED

If Light Path diagnostics points to PCI LED, reseat the PCI boards.

Myrinet communication fails

If communication fails over the Myrinet switch then check the following:

- If the Myricom adapter card green LED light is not on, check the cable connector for correct polarity (transmit/receive).
- Check to see that the GM module is installed by running the `lsmod` command.
- Check to see if the Myricom adapter is up and running by using the `ifconfig` command.

Appendix E. Configuring network switches

General networking notes

When setting up switches in the 512 mode or any time there intentionally are multiple connections between switches you must designate one of the core switches as the spanning tree root. In the case of the 512 node configuration it must be one of the Cisco Gigabit 4006 switches.

When setting up VLANs on a Cisco Gigabit 4006 running the Cisco Catalyst operating system (CatOS) make sure to set the vtp domain name. This can be any name since we are not using vtp to maintain the VLANs.

The Cisco Ethernet 3508/3524 switches only have 1 virtual Ethernet port. This can be assigned to any VLAN on the switch. Which ever VLAN it is assigned to should be designated as the Management VLAN for the switch.

The Cisco Gigabit 4006 switch running IOS can have an IP connection for each VLAN. However the management port on the Supervisor card can only be used for recovery situations. The number 1 port in the Management VLAN can be dedicated to hook up the management network to the Cisco Gigabit 4006 switch.

The Cisco Gigabit 4006 switch running the Cisco Catalyst Operating System (CatOS) has one port that can be used as an Ethernet connection. It is the **sc0** port and can be utilized in any VLAN. The sc0 port must be assigned to the Management VLAN. Again one port assigned to the Management VLAN needs to be reserved to make the connection to the switch itself.

Load balancing across EtherChannels is an important performance point. This is something that would be unique to the jobs that the customer intends to run on the cluster.

To split networks, creating a primary cluster VLAN and a Management VLAN in the switches, requires an extra connection between the 3550 and the 3508.

The RJ45 (copper) adapter GBICS must be connected to a Gigabit port or the link fails. Those GBICS will not negotiate speed.

The Linux kernel by default supports proxy arping. This can cause problems on a shared media network. If you have more that one NIC in the same broadcast domain there is a known problem with proxy arping. Proxy arping allows either interface in the broadcast domain to respond to an arp request. This can cause IP traffic to be handled by an interface other than the intended one. The only way to prevent this is to create separate VLANs in the switches.

Switch commands

Use the following section for information on switch commands. Use the command line interface for executing commands.

Switch commands for the Cisco Catalyst Ethernet 3550 switch running IOS

These commands will work with a Cisco Gigabit 4006 switch running the Cisco IOS software platform as well. To set up VLANs, at the command prompt, type:

© Copyright IBM Corp. 2004

```
vlan database
vtp transparent
vlan <id>  name <string>
exit
```

To assign ports to the VLAN, type:

```
conf t
int mod/port
switchport access vlan <id>
end
```

To set Ethernet address for switch assign to Management VLAN, type:

```
conf t
int vlan <id>
ip address <ip address> <netmask>
managment
end
```

To assign a name to the switch, type:

```
set system name <some string>
conf t
hostname <string>
end
```

To see the VLAN setup, type:

```
show vlan
```

To see spanning-tree protocol information on a port-by-port basis, type:

```
show spanning-tree brief //
```

Switch commands for the Cisco Gigabit 4006 switch running IOS

The following commands also work with the Cisco 3550 Ethernet switch running IOS. To set up VLANs, at the command prompt, type:

```
conf t
vlan <id>
name <management network>
end
```

To assign ports to the VLAN, type:

```
conf t
vlan <id>
name <management network>
end
```

To set Ethernet address for switch assigned to Management VLAN, type:

```
conf t
int vlan <id>
ip address <ip address> <netmask>
end
```

To assign a name to the switch, type:

```
set system name <some string>
conf t
hostname <string>
end
```

To create an EtherChannel, type:

```
conf t
int range <mode/port> - <port>
channel-group <id> mode desirable non-silent
end
```

To remove an EtherChannel, type:

```
conf t
int range <mode/port> - <port>
no channel-group
end
```

For the following command to work make sure all ports in the group are set up identically.

```
conf t
int range <mode/port> - <port>
channel-group <id> mode desirable on
end
```

This command generates an error message about port differences. Once the command completes, set it back to the desirable mode.

To turn off the spanning-tree protocol on ports going to cluster and storage nodes, type:

```
conf t
int range <mode/port> - <port>
switchport host
end
```

To see the VLAN setup, type:

```
show vlan
```

To set the switch as the spanning-tree protocol root. Run the command once for each VLAN:

```
conf t
spanning-tree <id> root primary
end
```

To set the switch as the spanning-tree protocol root secondary. Run the command once for each VLAN:

```
conf t
spanning-tree <id> root secondary
end
```

See spanning-tree protocol root information on a port-by-port basis:

```
show spanning-tree brief
```

See EtherChannels that are up and running:

```
show etherchannel
```

Switch commands for the Cisco Gigabit 4006 switch running Cisco Catalyst Operating System

To set up VLANs, type:

```
set vtp domain <string>
set vlan <2> name <management-network>
```

To assign ports to the VLAN, type:

```
set vlan <2> 2/1-10
```

To set Ethernet address for switch assigned to Management VLAN, type:

```
set interface sc0 <2> <172.30.50.3/255.255.0.0>
```

To set switch interface to a VLAN, type:

```
set interface sc0 <2>
```

To create an EtherChannel, type:

```
set port channel mod/port mode desirable non-silent
```

To assign a name to the switch, type:

```
set system name <some string>
```

For the following command to succeed make sure all ports in the group are set up identically. If an EtherChannel does not form, type:

```
set port channel <mod/port> mode on
```

This command generates an error message about port differences. Once the command succeeds, set it back to the desired mode.

To turn off the spanning-tree protocol on ports going to compute and storage nodes, type:

```
set port host
```

To see the VLAN setup, type:

```
show vlan
```

To set the switch as the spanning-tree protocol primary, type this command once for each VLAN:

```
set spantree root <vlanid>
```

To set the switch as the spanning-tree protocol secondary, type this command once for each VLAN:

```
set spantree root secondary <vlanid>
```

To view spanning-tree protocol root information on a port-by-port basis, type:

```
show spantree
```

To view EtherChannels that are up and running, type:

```
show channel
```

To disable and re-enable an EtherChannel that fails to link up, type:

```
set port disable <mod/port>
set port enable <mod/port>
```

Miscellaneous CISCO switch commands for the Cisco Catalyst Operating System

To clear configuration information from all modules in the switch, type:

```
clear config <all>
```

To clear configuration information from a module, type:

```
clear config <mod>
```

To view what ports are blocked by the spanning-tree protocol, type:

```
show spantree
```

Miscellaneous Cisco switch commands for IOS

To view the ports that are blocked by the spanning-tree protocol, type:

```
show sp br
```

Appendix F. Notices

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

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

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

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

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

Any references in this information to non-IBM Web sites 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 IBM product, and use of those Web sites is at your own risk.

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

Edition notice

© Copyright International Business Machines Corporation 2004. All rights reserved.

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

Trademarks

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

Active Memory
Active PCI
Active PCI-X
Alert on LAN
BladeCenter
C2T Interconnect
Chipkill
EtherJet
e-business logo
@server
FlashCopy
IBM
IBM (logo)
IntelliStation
NetBAY
Netfinity
NetView
OS/2 WARP
Predictive Failure Analysis
PS/2
ServeRAID
ServerGuide
ServerProven
TechConnect
ThinkPad
Tivoli
Tivoli Enterprise
TotalStorage
Update Connector
Wake on LAN
XA-32
XA-64
X-Architecture
XceL4
XpandOnDemand
xSeries

Lotus, Lotus Notes, SmartSuite, and Domino are trademarks of Lotus Development Corporation and/or IBM Corporation in the United States, other countries, or both.

Intel, MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

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

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

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

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

Important notes

Processor speeds indicate the internal clock speed of the microprocessor; other factors also affect application performance.

CD-ROM drive speeds list the variable read rate. Actual speeds vary and are often less than the maximum possible.

When referring to processor storage, real and virtual storage, or channel volume, KB stands for approximately 1000 bytes, MB stands for approximately 1 000 000 bytes, and GB stands for approximately 1 000 000 000 bytes.

When referring to hard disk drive capacity or communications volume, MB stands for 1 000 000 bytes, and GB stands for 1 000 000 000 bytes. Total user-accessible capacity may vary depending on operating environments.

Maximum internal hard disk drive capacities assume the replacement of any standard hard disk drives and population of all hard disk drive bays with the largest currently supported drives available from IBM.

Maximum memory may require replacement of the standard memory with an optional memory module.

IBM makes no representation or warranties regarding non-IBM products and services that are ServerProven®, including but not limited to the implied warranties of merchantability and fitness for a particular purpose. These products are offered and warranted solely by third parties.

IBM makes no representations or warranties with respect to non-IBM products. Support (if any) for the non-IBM products is provided by the third party, not IBM.

Some software may differ from its retail version (if available), and may not include user manuals or all program functionality.

Product recycling and disposal

This unit contains materials such as circuit boards, cables, electromagnetic compatibility gaskets, and connectors which may contain lead and copper/beryllium alloys that require special handling and disposal at end of life. Before this unit is disposed of, these materials must be removed and recycled or discarded according to applicable regulations. IBM offers product-return programs in several countries. Information on product recycling offerings can be found on IBM's Internet site at http://www.ibm.com/ibm/environment/products/prp.shtml.

Battery return program

This product may contain a 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 collection process for reuse, recycling, or proper disposal of used IBM sealed lead acid, nickel cadmium, nickel metal hydride, and battery packs from IBM equipment. For information on proper disposal of these batteries, contact IBM at 1-800-426-4333. Have the IBM part number listed on the battery available prior to your call.

In the Netherlands, the following applies.

Electronic emission 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.

Industry Canada Class A emission compliance statement

This Class A digital apparatus complies with Canadian ICES-003.

Avis de conformité à la réglementation d'Industrie Canada

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

Australia and New Zealand Class A statement

Attention: 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.

United Kingdom telecommunications safety requirement

Notice to Customers

This apparatus is approved under approval number NS/G/1234/J/100003 for indirect connection to public telecommunication systems in the United Kingdom.

European Union EMC Directive conformance 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 nonrecommended 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 CISPR 22/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.

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 any interference caused by using other than recommended cables and connectors.

Attention: 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.

Taiwanese Class A warning statement

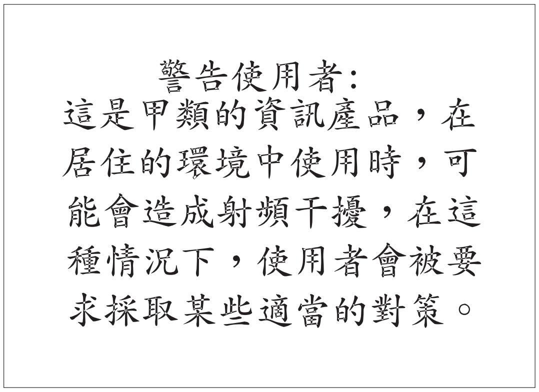

警告使用者:
這是甲類的資訊產品，在
居住的環境中使用時，可
能會造成射頻干擾，在這
種情況下，使用者會被要
求採取某些適當的對策。

Chinese Class A warning statement

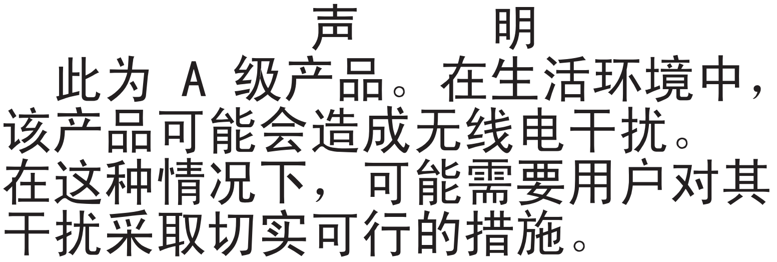

声　明
此为 A 级产品。在生活环境中，
该产品可能会造成无线电干扰。
在这种情况下，可能需要用户对其
干扰采取切实可行的措施。

Japanese Voluntary Control Council for Interference (VCCI) statement

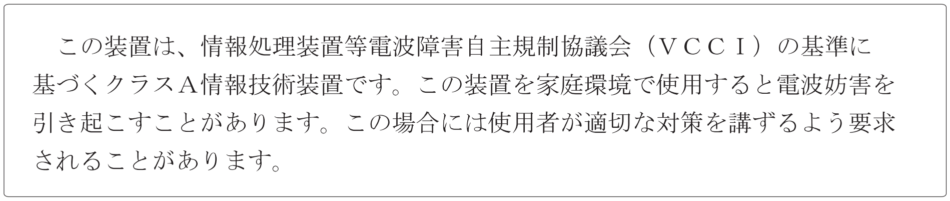

この装置は、情報処理装置等電波障害自主規制協議会（ＶＣＣＩ）の基準に基づくクラスＡ情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

International License Agreement for Non-Warranted Programs

Part 1 - General Terms

PLEASE READ THIS AGREEMENT CAREFULLY BEFORE USING THE PROGRAM. IBM WILL LICENSE THE PROGRAM TO YOU ONLY IF YOU FIRST ACCEPT THE TERMS OF THIS AGREEMENT. BY USING THE PROGRAM YOU AGREE TO THESE TERMS. IF YOU DO NOT AGREE TO THE TERMS OF THIS AGREEMENT, PROMPTLY RETURN THE UNUSED PROGRAM TO THE PARTY (EITHER IBM OR ITS RESELLER) FROM WHOM YOU ACQUIRED IT TO RECEIVE A REFUND OF THE AMOUNT YOU PAID.

The Program is owned by International Business Machines Corporation or one of its subsidiaries (IBM) or an IBM supplier, and is copyrighted and licensed, not sold.

The term ″Program″ means the original program and all whole or partial copies of it. A Program consists of machine-readable instructions, its components, data, audio-visual content (such as images, text, recordings, or pictures), and related licensed materials.

This Agreement includes **Part 1 - General Terms**, **Part 2 - Country-unique Terms**, and **License Information** and is the complete agreement regarding the use of this Program, and replaces any prior oral or written communications between you and IBM. The terms of **Part 2** and **License Information** may replace or modify those of **Part 1**.

1. **License**

   Use of the Program: IBM grants you a nonexclusive license to use the Program. You may 1) use the Program to the extent of authorizations you have acquired and 2) make and install copies to support the level of use authorized, providing you reproduce the copyright notice and any other legends of ownership on each copy, or partial copy, of the Program. If you acquire this Program as a program upgrade, your authorization to use the Program from which you upgraded is terminated. You will make sure that anyone who uses the Program does so only in compliance with the terms of this Agreement. You may not 1) use, copy, modify, or distribute the Program except as provided in this Agreement; 2) reverse assemble, reverse compile, or otherwise translate the Program except as specifically permitted by law without the possibility of contractual waiver; or 3) sublicense, rent, or lease the Program. Transfer of Rights and Obligations You may transfer all your license rights and obligations under a Proof of Entitlement for the Program to another party by transferring the Proof of Entitlement and a copy of this Agreement and all documentation. The transfer of your license rights and obligations terminates your authorization to use the Program under the Proof of Entitlement.

2. **Proof of Entitlement**

   The Proof of Entitlement for this Program is evidence of your authorization to use this Program and of your eligibility for any future upgrade program prices (if announced), and potential special or promotional opportunities.

3. **Charges and Taxes**

   IBM defines use for the Program for charging purposes and specifies it in the Proof of Entitlement. Charges are based on extent of use authorized. If you wish to increase the extent of use, notify IBM or its reseller and pay any applicable charges. IBM does not give refunds or credits for charges already due or paid.

© Copyright IBM Corp. 2004

If any authority imposes a duty, tax, levy or fee, excluding those based on IBM's net income, upon the Program supplied by IBM under this Agreement, then you agree to pay that amount as IBM specifies or supply exemption documentation.

4. **No Warranty**

   SUBJECT TO ANY STATUTORY WARRANTIES WHICH CAN NOT BE EXCLUDED, IBM MAKES NO WARRANTIES OR CONDITIONS EITHER EXPRESS OR IMPLIED, INCLUDING WITHOUT LIMITATION, THE WARRANTY OF NON-INFRINGEMENT AND THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, REGARDING THE PROGRAM OR TECHNICAL SUPPORT, IF ANY. IBM MAKES NO WARRANTY REGARDING THE CAPABILITY OF THE PROGRAM TO CORRECTLY PROCESS, PROVIDE AND/OR RECEIVE DATE DATA WITHIN AND BETWEEN THE 20TH AND 21ST CENTURIES.

   The exclusion also applies to any of IBM's subcontractors, suppliers, or program developers (collectively called ″Suppliers″).

   Manufacturers, suppliers, or publishers of non-IBM Programs may provide their own warranties.

5. **Limitation of Liability**

   NEITHER IBM NOR ITS SUPPLIERS WILL BE LIABLE FOR ANY DIRECT OR INDIRECT DAMAGES, INCLUDING WITHOUT LIMITATION, LOST PROFITS, LOST SAVINGS, OR ANY INCIDENTAL, SPECIAL, OR OTHER ECONOMIC CONSEQUENTIAL DAMAGES, EVEN IF IBM IS INFORMED OF THEIR POSSIBILITY. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THE ABOVE EXCLUSION OR LIMITATION MAY NOT APPLY TO YOU.

6. **General**

   Nothing in this Agreement affects any statutory rights of consumers that cannot be waived or limited by contract.

   IBM may terminate your license if you fail to comply with the terms of this Agreement. If IBM does so, your authorization to use the Program is also terminated and you must immediately destroy the Program and all copies you made of it.

   You agree to comply with applicable export laws and regulations.

   Neither you nor IBM will bring a legal action under this Agreement more than two years after the cause of action arose unless otherwise provided by local law without the possibility of contractual waiver or limitation.

   Neither you nor IBM is responsible for failure to fulfill any obligations due to causes beyond its control. The laws of the country in which you acquire the Program govern this Agreement, except 1) in Australia, the laws of the State or Territory in which the transaction is performed govern this Agreement; 2) in Albania, Armenia, Belarus, Bosnia/Herzegovina, Bulgaria, Croatia, Czech Republic, Federal Republic of Yugoslavia, Georgia, Hungary, Kazakhstan, Kirghizia, Former Yugoslav Republic of Macedonia (FYROM), Moldova, Poland, Romania, Russia, Slovak Republic, Slovenia, and Ukraine, the laws of Austria govern this Agreement; 3) in the United Kingdom, all disputes relating to this Agreement will be governed by English Law and will be submitted to the exclusive jurisdiction of the English courts; 4) in Canada, the laws in the Province of Ontario govern this Agreement; and 5) in the United States and Puerto Rico, and People's Republic of China, the laws of the State of New York govern this Agreement.

Part 2 - Country-unique Terms

AUSTRALIA: No Warranty (Section 4): The following paragraph is added to this Section: Although IBM specifies that there are no warranties, you may have certain rights under the Trade Practices Act 1974 or other legislation and are only limited to the extent permitted by the applicable legislation.

Limitation of Liability (Section 5): The following paragraph is added to this Section: Where IBM is in breach of a condition or warranty implied by the Trade Practices Act 1974, IBM's liability is limited to the repair or replacement of the goods, or the supply of equivalent goods. Where that condition or warranty relates to right to sell, quiet possession or clear title, or the goods are of a kind ordinarily acquired for personal, domestic or household use or consumption, then none of the limitations in this paragraph apply.

GERMANY: No Warranty (Section 4): The following paragraphs are added to this Section: The minimum warranty period for Programs is six months. In case a Program is delivered without Specifications, we will only warrant that the Program information correctly describes the Program and that the Program can be used according to the Program information. You have to check the usability according to the Program information within the ″money-back guarantee″ period.

Limitation of Liability (Section 5): The following paragraph is added to this Section: The limitations and exclusions specified in the Agreement will not apply to damages caused by IBM with fraud or gross negligence, and for express warranty.

INDIA: General (Section 6): The following replaces the fourth paragraph of this Section: If no suit or other legal action is brought, within two years after the cause of action arose, in respect of any claim that either party may have against the other, the rights of the concerned party in respect of such claim will be forfeited and the other party will stand released from its obligations in respect of such claim.

IRELAND: No Warranty (Section 4): The following paragraph is added to this Section: Except as expressly provided in these terms and conditions, all statutory conditions, including all warranties implied, but without prejudice to the generality of the foregoing, all warranties implied by the Sale of Goods Act 1893 or the Sale of Goods and Supply of Services Act 1980 are hereby excluded.

ITALY: Limitation of Liability (Section 5): This Section is replaced by the following: Unless otherwise provided by mandatory law, IBM is not liable for any damages which might arise.

NEW ZEALAND: No Warranty (Section 4): The following paragraph is added to this Section: Although IBM specifies that there are no warranties, you may have certain rights under the Consumer Guarantees Act 1993 or other legislation which cannot be excluded or limited. The Consumer Guarantees Act 1993 will not apply in respect of any goods or services which IBM provides, if you require the goods or services for the purposes of a business as defined in that Act.

Limitation of Liability (Section 5): The following paragraph is added to this Section: Where Programs are not acquired for the purposes of a business as defined in the Consumer Guarantees Act 1993, the limitations in this Section are subject to the limitations in that Act.

PEOPLE'S REPUBLIC OF CHINA: Charges (Section 3): The following paragraph is added to the Section: All banking charges incurred in the People's Republic of China will be borne by you and those incurred outside the People's Republic of China will be borne by IBM.

UNITED KINGDOM: Limitation of Liability (Section 5): The following paragraph is added to this Section at the end of the first paragraph: The limitation of liability will not apply to any breach of IBM's obligations implied by Section 12 of the Sale of Goods Act 1979 or Section 2 of the Supply of Goods and Services Act 1982.

License Information

Program-unique Terms

The following terms and conditions are in addition to those of the IBM International License Agreement for Non-Warranted Programs (ILAN). Solely with respect to your use of the Cisco software (the ″Cisco Software″) contained within the IBM product you have purchased.

1. Your license to the Cisco Software is a license to (a) use the software in the operation of a Cisco networking product only; (b) make not more than one (1) copy of the Cisco Software, which you may use only for purposes of backup and disaster recovery. You may not otherwise copy the Cisco Software, and you may not transfer the Cisco Software, even if you sell or lease the Cisco networking product with which the Cisco Software is provided. The purchaser or other transferee of the Cisco Software must obtain from Cisco or a Cisco reseller (including IBM) a new license to use the Cisco Software.
2. In addition to the warranty disclaimers provided in Point 4 of the ILA, Cisco disclaims any warranty that the Cisco Software or any equipment, system or network on which the Cisco Software is used will be free of vulnerability to intrusion or attack.
3. In the event you breach any provision of the ILA provided to you, or any provision of these additional terms, your right to use the Cisco Software will terminate immediately.
4. If you received the Cisco Software in the European Union, the Middle East, or Africa, the law applicable to your use of the Cisco Software is English law. If you received the Cisco Software in Canada, the law applicable to your use of the Cisco Software is Ontario law. If you received the Cisco Software in Australia or New Zealand, the law applicable to your use of the Cisco Software is Australian law. If you received the Cisco Software elsewhere in the world, the law applicable to your use of the Cisco Software is the law of the State of California, the United States of America.
5. For United States government users, the Cisco Software is Commercial Computer Software provided with Restricted Rights per the terms of the Federal Acquisition Regulation.
6. In the event you receive upgrades to the Cisco Software, you may only use such upgrades if, at the time you receive them, you have a valid license to use the Cisco Software which was upgraded or updated.

Index

© Copyright IBM Corp. 2004

Part Number: 25K8420

Printed in USA

(1P) P/N: 25K8420