




IBM @server Cluster 1350
Installation and Service Guide





@server

IBM @server Cluster 1350
Installation and Service Guide

Note: Before using this information and the product it supports, read the general information in “Safety” on page vii and Appendix F, “Notices,” on page 67.

Eighth edition, July 2005

© Copyright International Business Machines Corporation 2005. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Safety	vii
Safety Information	vii
Handling static-sensitive devices	xv
Notices and statements used in this document	xv
Chapter 1. System overview	1
Cluster components	4
Cluster nodes	5
Management node	5
Storage nodes	5
Storage servers	6
Storage expansion units	6
Fibre Channel adapters, switches, and expansion cards	6
Hard disk drives	6
ServeRAID storage controller adapters	7
Console	7
KVM switches	7
InfiniBand switches and devices	7
InfiniBand host channel adapters	7
10/100/1000 Ethernet switches	8
10 Gigabit Ethernet switches	8
I/O modules	8
Terminal servers	8
High-speed Myrinet switches	8
Myrinet BladeCenter connectivity	9
High-speed Cisco switches and line cards	9
Power distribution unit (PDU)	9
Related documentation	10
Chapter 2. Cabinet placement and intercabinet cabling	11
Minimum service clearances	11
Types of intercabinet cabling	12
Cabinet placement guidelines	13
Installing a stability kit	15
Engaging the side stabilizers	15
Floor loading	15
Preparing a Rear Door Heat eXchanger	16
Chapter 3. Cabling the Cluster 1350	17
VLAN options	18
Connecting the cables	24
1 Gigabit Ethernet cabling	25
Myrinet high-speed switch cabling	25
10/100/1000 Ethernet cabling	25
Fibre-channel cabling	25
KVM cabling	25
Remote console manager cabling	26
Local console manager cabling	26
Topspin InfiniBand switch and host channel adapter cabling	26
Voltaire InfiniBand switch router cabling	26
Replacing a defective cable in a harness	26
Chapter 4. Checking the Cluster 1350 cabling process	29

Installing cables and checking connections in the expansion cabinets	29
Installing cables and checking connections in the primary cabinet	30
Turning on the power to the expansion cabinets.	30
Turning on the power to the primary cabinet	30
Verifying the hardware installation	31
Diagnosing a lights out or brown out event.	32
Related topics	33
Chapter 5. Accessing the cluster from a remote location	35
Using the remote power command	35
Remote console	35
Related topics	35
Chapter 6. Troubleshooting hardware and software problems	37
Displaying node configuration information	37
Isolating network, node, and Linux problems	37
Clustering with one network	38
Clustering with two networks	39
Isolating hardware problems	42
Node checks.	42
Service processor network	44
Checking storage	45
Checking the terminal server.	46
Troubleshooting the KVM network	47
File-system failure.	47
PFA alert indicates internal disk.	47
I/O errors in syslog	47
Isolating software problems	48
Operating-system checks	48
CSM checks	48
GPFS checks	49
SNMP monitoring	49
Configuring SNMP alerts from Myrinet	49
Resetting the Remote Supervisor Adapter	50
Checking service processor logs	50
Management, cluster, and storage node problems	51
Disk drive failures	51
System board failures	52
xSeries 336 problems	53
BladeCenter problems	53
Power problems	53
Related documentation	54
Appendix A. Getting help and technical assistance	55
Before you call	55
Using the documentation	55
Getting help and information from the World Wide Web	56
Software service and support	56
Hardware service and support	56
Appendix B. Frequently asked questions	57
Appendix C. Error and event logs	59
Appendix D. Known problems	61
Node	61

Amber light on node	61
COM port settings in BIOS	61
CSM.	61
Stale NFS mounts.	61
rpower hard shut down	61
Storage	61
Driver module ordering	61
KVM.	62
GUI does not appear on first node.	62
2x8 Switch powers on with console port B	62
Cluster port 1 reboots	62
Subsequent KVMs unresponsive	62
RSA and service processor	63
RSA unable to load firmware	63
RSA or service processor invalid naming	63
Light path points to PCI LED.	63
Myrinet communication fails	63
Appendix E. General networking notes	65
Appendix F. Notices	67
Edition notice	67
Trademarks	68
Important notes.	68
Product recycling and disposal	69
Battery return program	70
Electronic emission notices	71
Federal Communications Commission (FCC) statement	71
Industry Canada Class A emission compliance statement	71
Australia and New Zealand Class A statement	71
United Kingdom telecommunications safety requirement.	71
European Union EMC Directive conformance statement.	72
Taiwanese Class A warning statement	72
Chinese Class A warning statement	72
Japanese Voluntary Control Council for Interference (VCCI) statement	72
Power cords	73
Appendix G. Declaration of IBM product noise emission values	75
International License Agreement for Non-Warranted Programs.	77
Part 1 - General Terms	77
Part 2 - Country-unique Terms	79
License Information	80
Index	81

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.

Safety Information

Before installing this product, read the Safety Information.

قبل تركيب هذا المنتج، يجب قراءة الملاحظات الأمنية

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.

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 pečítajte 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

Note:

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 Safety Information* book.

For example, if a caution statement begins with a number 1, translations for that caution statement appear in the *IBM 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 1:



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.**

Statement 3:



CAUTION:

When laser products (such as CD-ROMs, DVD drives, fiber optic devices, or transmitters) are installed, note the following:

- **Do not remove the covers. Removing the covers of the laser product could result in exposure to hazardous laser radiation. There are no serviceable parts inside the device.**
- **Use of controls or adjustments or performance of procedures other than those specified herein might result in hazardous radiation exposure.**



DANGER

Some laser products contain an embedded Class 3A or Class 3B laser diode. Note the following.

Laser radiation when open. Do not stare into the beam, do not view directly with optical instruments, and avoid direct exposure to the beam.




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

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:



CAUTION:

Never remove the cover on a power supply or any part that has the following label attached.



Hazardous voltage, current, and energy levels are present inside any component that has this label attached. There are no serviceable parts inside these components. If you suspect a problem with one of these parts, contact a service technician.

Statement 9:



CAUTION:

To avoid personal injury, disconnect the hot-swap fan cables before removing the fan from the device.

Statement 13:



DANGER

Overloading a branch circuit is potentially a fire hazard and a shock hazard under certain conditions. To avoid these hazards, ensure that your system electrical requirements do not exceed branch circuit protection requirements. Refer to the information that is provided with your device for electrical specifications.

Statement 14:



CAUTION:

Hazardous voltage, current, and energy levels might be present. Only a qualified service technician is authorized to remove the covers where the following label is attached.



Statement 15:



CAUTION:

Make sure that the rack is secured properly to avoid tipping when the server unit is extended.

Statement 16:



CAUTION:


Some accessory or option board outputs exceed Class 2 or limited power source limits and must be installed with appropriate interconnecting cabling in accordance with the national electric code.

Statement 19:



CAUTION:

The power-control button on the device does not turn off the electrical current supplied to the device. The device also might have more than one connection to dc power. To remove all electrical current from the device, ensure that all connections to dc power are disconnected at the dc power input terminals.



Statement 20:



CAUTION:

To avoid personal injury, before lifting the unit, remove all the blades to reduce the weight.

Statement 21:



CAUTION:

Hazardous energy is present when the blade is connected to the power source. Always replace the blade cover before installing the blade.

Statement 23:



CAUTION:

Do not place any object weighing more than 50 kg (110 lb.) on top of rack-mounted devices.



>50 kg (110 lb.)



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.



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.
- 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.

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 or rack 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 server 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 at <http://www.ibm.com/pc/support/>. 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 an overview of the IBM Cluster 1350 software components, hardware components, and related documentation.

The Cluster 1350 identifies two types of cabinet: primary and expansion. A cabinet is called primary if it contains the management node and console monitor. An expansion cabinet can contain storage nodes or mass-storage devices as well as computing nodes, called cluster nodes; it does not contain a management node or a console.

The cluster consists of:

- One or more racks.
- From 4 to 1025 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[®] 346 or @server[™] 326) for cluster systems management (CSM) 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 switches.
- A cluster VLAN used for other management traffic and user traffic. 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 (such as, the xSeries 336) and Remote Supervisor Adapter (RSA) devices. The first node in a serial connection must have an RSA 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, an additional Gigabit Ethernet, or an InfiniBand switch.
- The customer can elect to configure a subset of cluster nodes with additional external storage. This can be a Fibre Channel solution.
- A supported distribution of the Linux operating system.
- Cluster system management software, such as, CSM.


The Cluster 1350 supports a maximum of 1025 nodes in addition to the one required xSeries 346 management node (or @server 326 management node for 64-bit processing environments). All nodes must run one of the following Linux versions:

- Red Hat Enterprise Linux (RHEL), version 3, for AMD or EMT (32-bit or 64-bit)
- RHEL, version 4, for AMD or EMT (32-bit or 64-bit)
- SUSE LINUX Enterprise Server (SLES), version 9, for AMD or EMT (32-bit or 64-bit)
- Linux Cluster Installation Tool (LCIT), version 3.4

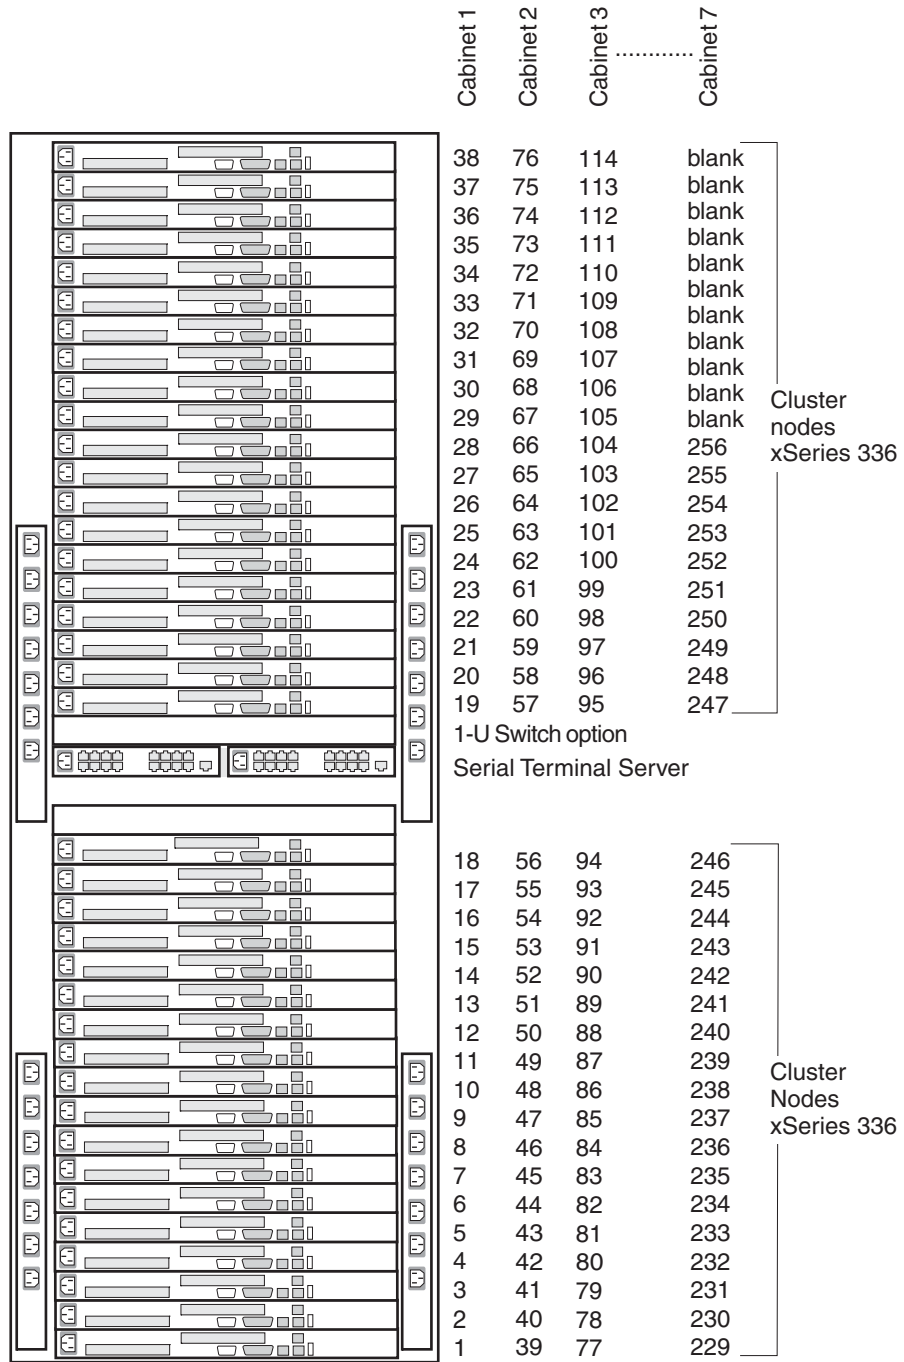
The Cluster 1350 uses a primary cabinet and an expansion cabinet. The primary cabinet contains the management node and console monitor. An expansion cabinet does not contain a management node or console. An expansion cabinet can contain the following components:

- Cluster or compute nodes
- Storage nodes
- Mass storage devices
- Expansion units
- Networking switches

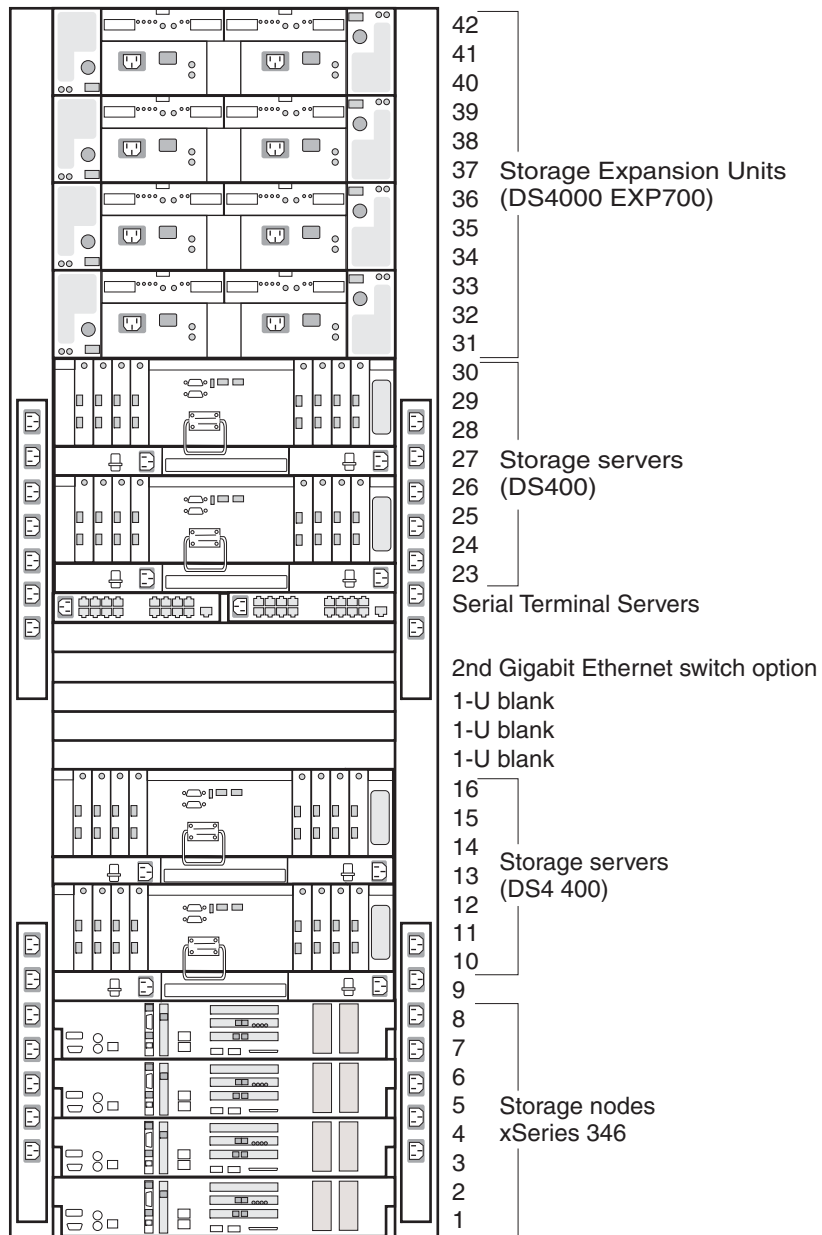
The following figures illustrate a primary cabinet, an expansion cabinet containing cluster nodes, and an expansion cabinet containing storage controllers and mass storage.



The following figure shows how the node numbering scheme maps to other expansion cabinets.



The following figure shows an 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
- Fibre Channel adapters, switches, and expansion cards
- Hard disk drives

- ServeRAID storage controller adapters
- Console
- Keyboard/video/mouse (KVM) switches
- InfiniBand switches and devices
- InfiniBand host channel adapters
- Fibre Channel adapters, switches, and expansion cards
- 10/100/1000 Ethernet switches, Force 10, and SAN switches
- I/O modules
- Terminal servers
- High-speed Myrinet switches
- High-speed Cisco switches and line cards
- 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 326
- xSeries 336
- xSeries 346
- @server BladeCenter™ populated with HS20 Type 8843 blade servers
- @server BladeCenter populated with JS20 Type 8842 blade servers
- @server BladeCenter populated with LS20 Type 8850 blade servers

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 346 server running Linux. You can also use an @server 326 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 326
- xSeries336
- xSeries 346
- @server BladeCenter unit populated with HS20 Type 8843 blade servers
- @server BladeCenter unit populated with JS20 Type 8842 blade servers
- @server BladeCenter unit populated with LS20 Type 8850 blade servers

Storage servers

For the storage server option, the Cluster 1350 system uses any of the following storage servers:

- IBM TotalStorage® DS4 300 storage server
- IBM TotalStorage DS4 300 with Turbo storage server
- IBM TotalStorage DS4 400 storage server
- IBM TotalStorage DS4 500 storage server
- IBM TotalStorage DS300 Storage Area Network (SAN) storage server
- IBM TotalStorage DS400 SAN storage server

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:

- IBM TotalStorage DS4000 EXP100 R2: supports DS4000 250 GB S-ATA hard disk drives
- IBM TotalStorage DS4000 EXP400: supports Ultra320 SCSI hard disk drives
- IBM TotalStorage DS4000 EXP700: supports Ultra320 SCSI hard disk drives
- IBM TotalStorage DS4300 EXP700

Fibre Channel adapters, switches, and expansion cards

The cluster supports the following Fibre Channel devices:

- Fibre Channel SCSI/iSCSI single channel adapter
- Fibre Channel SCSI/iSCSI dual channel adapter
- Brocade Fibre Channel switch
- IBM @server BladeCenter Fibre Channel switch module
- IBM @server BladeCenter Fibre Channel expansion card
- QLogic Quad Channel 2 Gb Fibre Channel host bus adapter
- Fibre Channel switch 2005 model H-16 (16-port)

Hard disk drives

The cluster supports the following hard disk drives:

- IBM DS4000 EXP100 250 GB S-ATA hard disk drive
- IBM DS4000 Ultra320 300 GB 10Khz SCSI hard disk drive
- IBM DS4000 Ultra320 146 GB 15Khz SCSI hard disk drive
- Maxtor 250 GB dual-port S-ATA hard disk drive.
- Savvio 2.5 inch small-form factor hard disk drive

ServeRAID storage controller adapters

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

- ServeRAID™ -6M Ultra 320 SCSI controller adapter
- ServeRAID 7t Serial-ATA controller adapter
- ServeRAID 7k Ultra 320 SCSI controller adapter

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 switches

The 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 @server 326 server.

InfiniBand switches and devices

The following InfiniBand switches and line cards are supported:

- Topspin 120 InfiniBand switch
- Topspin 270 InfiniBand switch (96-port)
 - Topspin 12-port 4X InfiniBand line card
 - Topspin 270 switch fabric module
 - Topspin InfiniBand management module
- Voltaire 9024 InfiniBand switch router (24-port)
- Voltaire 9288 switch bundle:
 - Chassis unit
 - Switch line power card
 - Power supply
 - Switch fabric

InfiniBand host channel adapters

The following InfiniBand host channel adapters are supported:

- Topspin InfiniBand host channel adapter PCI-X
- Topspin InfiniBand host channel adapter PCI-Express
- Voltaire InfiniBand host channel adapter PCI-X
- Voltaire InfiniBand host channel adapter PCI-Express

10/100/1000 Ethernet switches

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

- IBM @server BladeCenter Gigabit Ethernet switch module (4-port)
- Cisco Catalyst Ethernet switch model 3750G-24T stackable Gigabit Ethernet switch (24-port)
- SMC Ethernet switch model 8624T (24-port)
- SMC Ethernet switch model 8648T (48-port)
- SMC 8748M Tiger IV switch (48-port - 10 Gb uplinks)
- Force 10 1U stackable switch (48-port - 2 x 10 Gb uplinks)
- Force 10 E600 switch chassis bundle

10 Gigabit Ethernet switches

The Cluster 1350 uses the following 10 Gb Ethernet switches:

- SMC 8708L2 Tiger IV switch (8-port)
- SMC 8748M Tiger IV XENPAK receiver tray

I/O modules

I/O modules provide an unconfigured network connection that enables the blade servers to connect to an existing network infrastructure. No configuration is required. The cluster supports the following I/O modules:

- IBM BladeCenter Gigabit Ethernet switch module (4-port)
- Copper Pass-thru module for IBM @server BladeCenter
- Cisco Gigabit Ethernet Switch for IBM @server BladeCenter
- Brocade Enterprise SAN Switch Module for IBM @server BladeCenter
- Brocade Entry SAN switch module for IBM @server BladeCenter
- IBM switch module for blades
- IBM BladeCenter Fibre Channel switch module (2-port)
- IBM Layer 2/Layer 3 Copper Gigabit Ethernet switch module

Terminal servers

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

- In-Reach LX-4032 switch (32-port)
- In-Reach LX-4048 switch(48-port)
- Cyclades ATP1265 switch (48-port)

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 or 4 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-PCIXF-2 line card
- M3F2-PCIXE-2 (low-profile PCI-X dual-port card)
- M3-SW16-8E (Ethernet line card)
- M3-CLOS-ENCL 14U, 20-slot enclosure (256-ports)
- M3-SPINE-ENCL 14U, 20-slot spine enclosure
- M3-4SW32-16Q line card (16-ports quad)
- M3-THRU-16Q (spine thru module)
- M3-SW32-16F line card (SW32 with 16 fiber port front connections)
- M3-2SW32 line card
- M3S-PCIXD-2-1 line card (2Mb)
- M3S-PCIXF-4-1 line card (4 Mb)

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

Myrinet BladeCenter connectivity

Myrinet BladeCenter connectivity includes a Myrinet PCI card, the Optical Pass thru module, I/O expansion card, and special cable options.

High-speed Cisco switches and line cards

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 series 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
- 6500 - CEF720 10 Gb Ethernet module (4-port)
- XENPAK 10 GB-SR

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 PDU
- IBM Distributed Power Interconnect (DPI®) Rack PDU
- IBM DPI Front-End PDU
- IBM DPI High Density PDU

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 documentation

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 at <http://www.ibm.com/pc/support/>.

To remove and configure a cluster component, see the documentation that comes with the component.

To check for updated documentation, related documentation, and technical updates, complete the following steps:

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. Cabinet placement and intercabinet cabling

This chapter provides information about the final placement and intercabinet cabling of the cluster cabinets, how to install the frame stabilizer foot and outrigger to support each cabinet, and preparations for installing a Rear Door eXchanger Assembly unit.

Minimum service clearances

Several factors influence final cabinet placement. In addition to adequate cable length and staying within safe floor loading limits for your installation site, observe the minimum clearances required for service of the racks in the Cluster 1350. You might want to plan for extra space around the racks to allow movement of racks for service.

Figure 1 shows the clearances needed for servicing a rack in the Cluster 1350. Note that the service clearances for the sides of the rack are not required for operation, but only for servicing the PDUs.




Figure 1. Cabinet clearances required for service

When planning the installation of a large cluster and you want to arrange the racks in several rows, you must maintain a minimum aisle width of 36 inches as shown in Figure 2 on page 12. This will allow enough space for the front and back covers of the racks to open. An aisle width of 48 inches might be more appropriate if you want to make sure that you have enough room to move the rack to gain access to the sides for servicing the PDUs.

If you purchase a Myrinet CLOS 256 switch with your Cluster 1350, make sure to rotate the switch rack 180 degrees so that the rear of the rack unit is facing forward. This placement makes sure that there is proper air flow for the Myrinet device. See the *Installation* guide that came with your Myrinet CLOS 256 component for more information or go to <http://www.myri.com/scs/#documentation>.




Figure 2. Minimum aisle width for Cluster 1350 installations

Final cabinet placement determines what will qualify as adequate cable lengths. This chapter gives guidelines for intercabinet (rack-to-rack) cabling. Use these guidelines to determine the best cable lengths for your installation.

Remember to install the stability kit and engage the outriggers when the cabinets are placed in their final location. See “Installing a stability kit” on page 15 for information about installing the stability kit that comes with each cabinet. See “Engaging the side stabilizers” on page 15 for information about engaging the side stabilizers.

Types of intercabinet cabling

Cabinet placement is partly determined by the cables that run between cabinets. Both optical and copper cables are involved.

All intercabinet (rack-to-rack) cables and cable harnesses come in varying lengths up to 25 meters. There are no layout restrictions.

Color-coded CAT5E intercabinet Ethernet cables are available for intercabinet cabling of the cluster components. The color-coded cables operate exactly as the existing cluster intercabinet Ethernet cables with the added advantage of being able to group like-devices by a specific color. The cable colors and lengths include:

- Blue - 0.6 m, 1 m, 1.5 m, 2 m, 2.5 m, 3 m, 10 m, 25 m
- Green - 0.6 m, 1 m, 1.5 m, 2 m, 2.5 m, 3 m, 10 m, 25 m
- Yellow - 0.6 m, 1 m, 1.5 m, 2 m, 2.5 m, 3 m, 10 m, 25 m

Note: In some clusters, the FRU intercabinet cables can also be white or gray.

Use the following table for the available lengths for the different cables and cable harnesses.

Table 1. Available lengths for fibre-based Ethernet and Fibre Channel cables

Length	Part number	Note
1 m	36L9973	Fibre cable, SC to SC
5 m	03K9306	Fibre cable, SC to SC
25 m	03K9305	Fibre cable, SC to SC
1 m	19K1247	Fibre cable, LC to LC
3 m	31P6327	Fibre cable, LC to LC
5 m	19K1248	Fibre cable, LC to LC
10 m	31P6326	Fibre cable, LC to LC
25 m	19K1249	Fibre cable, LC to LC
0.208 m	19K1250	Fibre cable, SC to LC

Cabinet placement guidelines

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 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.
- 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.
- If you have purchased an IBM Rear Door Heat eXchanger, make sure that you have prepared the installation area with the proper fluid access, fluid connections, and fluid drainage. See “Preparing a Rear Door Heat eXchanger” on page 16 for more information.


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.

Installing a stability kit

A stability kit (IBM part number 31L8305), containing one stability kit, is included with the 1410 rack. Install the stability kit when the rack is at the final installation location. Figure 3 shows the attachment of the stability kit.

Figure 3. Cluster 1350 stability kit



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.

Engaging the side stabilizers

The rack cabinet provides side stabilizers, to give additional support to the rack unit. Engage the side stabilizers when the rack is at the final installation location. For more information about how to engage the stabilizers, see the documentation that came with your rack cabinet.

Floor loading

To determine the clearances required for each rack, you must know the maximum floor load rating in pounds per square foot or kilograms per square meter at your installation site. Table 2 on page 16 shows the necessary clearances or weight-distribution distances for a fully configured rack. The weight-distribution distances in all four directions indicate the area around the rack (minus covers) necessary to distribute the weight beyond the perimeter of the rack.

Weight-distribution areas must not overlap with adjacent computer equipment weight-distribution areas. You must also take into consideration the additional weight of the cables.

For example, if the raised floor at your site could safely support 351.5 kg/m² (72.0 lb./ft.²) then the correct distance between the sides of each cabinet would be 1524 mm (40 inches). You would also maintain front and back clearances of 762 mm (30 inches), making sure that no other equipment was placed within the area.

The Cluster 1350 cabinet footprint without covers is 622 mm (24.5 inches) by 1021 mm (40.2 inches). The values shown in the Table 2 assume a maximum rack weight of 1003 kilograms (2211 pounds). Floor loading will vary by configuration.

If you need to know the approximate point load that each caster exerts on a raised floor panel, divide the weight of the rack by four.

Table 2. Floor loading

lb./ft. ²	kg/m ²	Side clearances	Front and back clearances
72.0	351.5	508 mm (20 in.)	762 mm (30 in.)
78.4	382.8	508 mm (20 in.)	610 mm (24 in.)
86.8	423.8	508 mm (20 in.)	457 mm (18 in.)
93.1	454.6	254 mm (10 in.)	762 mm (30 in.)
102.4	500.0	254 mm (10 in.)	610 mm (24 in.)
114.5	559.1	254 mm (10 in.)	457 mm (18 in.)
112.8	550.8	127 mm (5 in.)	762 mm (30 in.)
124.8	609.3	127 mm (5 in.)	610 mm (24 in.)
140.5	686.0	127 mm (5 in.)	457 mm (18 in.)
139.3	680.1	25 mm (1 in.)	762 mm (30 in.)
154.9	756.3	25 mm (1 in.)	610 mm (24 in.)
175.4	856.4	25 mm (1 in.)	457 mm (18 in.)

Preparing a Rear Door Heat eXchanger

If you purchased an IBM Rear Door Heat eXchanger with your Cluster 1350, it is attached to the rack unit by an IBM service representative during the hardware installation.

After a door is installed, the customer is responsible for filling the fluid, draining the fluid, and connecting or disconnecting the plumbing connections when the unit arrives or is replaced by an IBM service representative. You might want to plan for extra space around the racks to allow movement of racks for fluid service. See the *Installation and Maintenance* guide that comes with the Rear Door Heat eXchanger or go to www.ibm.com/pc/support/ for more information.

Chapter 3. 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 RSAs. 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.

Primary cluster VLAN

The primary cluster VLAN provides a 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 10/100/1000 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 Myrinet switch (2 Gb or 4 Gb) for communication with cluster and storage nodes. The following options are available for the secondary cluster VLAN:

- 10/100/1000 Ethernet connection
- 2 Gb or 4 Gb Myrinet connection

Keyboard/video/mouse

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 RSA adapters 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 3. Type 1 10/100 Ethernet VLAN

Device	Management VLAN	10/100 primary cluster VLAN	Comments
Management node	Ethernet 2 connects to a Cisco 3750 G or an SMC switch		
KVM switch	Connects to a Cisco 3750 G or an SMC switch		
In-Reach LX-4000 terminal server	Connects to a Cisco 3750 G or an SMC switch		
APC switch	Connects to a Cisco 3750 G or an SMC switch		
Cluster nodes		Ethernet 0 connects to a Cisco 3750 G or an SMC switch or an SMC switch	
Storage nodes	Ethernet 2 connects to a Cisco 3750 G or an SMC switch		
DS4 300	Connects to a Cisco 3750 G or an SMC switch		Uses both jacks
DS4 400	Connects to a Cisco 3750 G or an SMC switch		Uses both jacks

Table 3. Type 1 10/100 Ethernet VLAN (continued)

Device	Management VLAN	10/100 primary cluster VLAN	Comments
DS4 500	Connects to a Cisco 3750 G or an SMC switch		Uses both jacks

Table 4. Type 2 10/100/1000 Ethernet VLAN

Device	Management VLAN	Gbit primary cluster VLAN	Comments
Management node	<ul style="list-style-type: none"> 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 4006 switch	<ul style="list-style-type: none"> Supervisor I connects to management-node Supervisor III connects to management-node 	<ul style="list-style-type: none"> 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	
DS4 300	Connects to Cisco 400x		Uses both jacks
DS4 400	Connects to Cisco 400x		Uses both jacks
DS4 500	Connects to Cisco 400x		Uses both jacks

Table 5. 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 3750 G or an SMC switch		
KVM switch	Connects to a Cisco 3750 G or an SMC switch		
In-Reach LX-4000 terminal server (32-port, 48-port)	Connects to a Cisco 3750 G or an SMC switch		

Table 5. Type 3 VLAN: 10/100 Ethernet with 10/100/1000 public high-speed VLAN (continued)

APC switch	Connects to a Cisco 3750 G or an SMC switch		
Cisco 4006 switch	<ul style="list-style-type: none"> Supervisor I connects to a Cisco 3750 G or an SMC switch Supervisor III connects to a Cisco 3750 G or an SMC switch 		
Cluster nodes		Ethernet 0 connects to a Cisco 3750 G or an SMC switch	Ethernet 2 connects to Cisco 400x Gbit
Storage nodes			Ethernet 2 connects to Cisco 400x Gbit
DS4 300	Connects to a Cisco 3750 G or an SMC switch		
DS4 400	Connects to a Cisco 3750 G or an SMC switch		
DS4 500	Connects to a Cisco 3750 G or an SMC switch		

Table 6. 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 3750 G		
KVM switch	Connects to a Cisco 3750 G or an SMC switch		
In-Reach LX-4000 terminal server (32-port, 48-port)	Connects to a Cisco 3750 G or an SMC switch		
APC switch	Connects to a Cisco 3750 G or an SMC switch		
Myrinet 32, 64, or 128 (both jacks), D card	Connects to a Cisco 3750 G or an SMC switch		
Topspin InfiniBand switch	Connects to a Cisco 3750 G or an SMC switch		
Voltaire 9024 InfiniBand switch	Connects to a Cisco 3750 G or an SMC switch		

Table 6. Type 4 VLAN: 10/100/1000 Ethernet with 2 Gbit public high-speed VLAN (continued)

Cluster nodes		Ethernet 0 connects to a Cisco 3750 G or an SMC switch	Myrinet adapter
Storage nodes	Ethernet 2 connects to a Cisco 3750 G or an SMC switch		Myrinet adapter
DS4 300	Connects to a Cisco 3750 G or an SMC switch		
DS4 400	Connects to a Cisco 3750 G or an SMC switch		
DS4 500	Connects to a Cisco 3750 G or an SMC switch		

Table 7. 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	<ul style="list-style-type: none"> Ethernet 1 Alias connects to Cisco 400x Ethernet 2 connects to Supervisor I or Supervisor III 	<ul style="list-style-type: none"> 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 4006 switch	<ul style="list-style-type: none"> Supervisor I connects to management-node Ethernet 2 Supervisor III connects to management-node Ethernet 2 	<ul style="list-style-type: none"> 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
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
DS4 400 (both jacks)	Connects to Cisco 400x		

Table 8. 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	<ul style="list-style-type: none"> Ethernet 1 Alias connects to Cisco 400x Ethernet 2 connects to Supervisor I or Supervisor III 	<ul style="list-style-type: none"> 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		
Voltaire 9024 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
DS4 400	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 9. Type 2 VLAN with multiple Cisco 400x switches

Device	Management VLAN	Gbit primary cluster VLAN	Comments
Management node	Ethernet 2 connects to a Cisco 3750 G or an SMC switch		
KVM switch	Connects to a Cisco 3750 G or an SMC switch		
In-Reach LX-4032 and 4048 (32-port, 48-port) terminal server	Connects a Cisco 3750 G or an SMC switch		

Table 9. Type 2 VLAN with multiple Cisco 400x switches (continued)

APC switch	Connects to a Cisco 3750 G or an SMC switch		
Cisco 4006 switch	<ul style="list-style-type: none"> • Supervisor I connects to a Cisco 3750 G or an SMC switch • Supervisor III connects to a Cisco 3750 G or an SMC switch 		
Cluster nodes		Ethernet 0 connects to Cisco 400x	
Storage nodes	Ethernet 2 connects to a Cisco 3750 G or an SMC switch	Ethernet 1 connects to Cisco 400x	
DS4 400	Connects to a Cisco 3750 G or an SMC switch		Uses both jacks

Table 10. 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 a Cisco 3750 G or an SMC switch		Fibre Channel PCI connects to public network
KVM switch	Connects to a Cisco 3750 G or an SMC switch		
In-Reach LX-4032 and 4048 (32-port, 48-port) terminal server	Connects to a Cisco 3750 G or an SMC switch		
APC switch	Connects to a Cisco 3750 G or an SMC switch		
Cisco 4006 switch	<ul style="list-style-type: none"> • Supervisor I connects to a Cisco 3750 G or an SMC switch • Supervisor III connects to a Cisco 3750 G or an SMC switch 		
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
DS4 400 (both jacks)	Connects to a Cisco 3750 G or an SMC switch		

Table 11. Type 6 VLAN with multiple Cisco 400x switches

Device	Management VLAN	Gbit primary cluster VLAN	Myrinet customer public high-speed VLAN
Management node	Ethernet 2 connects to Cisco 3750 G		Fibre Channel PCI adapter connects to public network
KVM switch	Connects to a Cisco 3750 G or an SMC switch		
In-Reach LX-4032 and 4048 (32-port, 48-port) terminal server	Connects to a Cisco 3750 G or an SMC switch		
APC switch	Connects to a Cisco 3750 G or an SMC switch		
Cisco 4006 switch	<ul style="list-style-type: none"> Supervisor I connects to a Cisco 3750 G or an SMC switch Supervisor III connects to a Cisco 3750 G or an SMC switch 		
Myrinet 32, 64, 128, or PCI adapter (both jacks)	Connects to a Cisco 3750 G or an SMC switch		
Topspin InfiniBand switch	Connects to a Cisco 3750 G or an SMC switch		
Voltaire 9024 InfiniBand switch router	Connects to Cisco 3750 G		
Cluster nodes		Ethernet 0 connects to Cisco 400x	Myrinet adapter
Storage nodes	Ethernet 2 connects to a Cisco 3750 G or an SMC switch	Ethernet 1 connects to Cisco 400x	Myrinet adapter
DS4 400	Connects to a Cisco 3750 G or an SMC switch		

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 or 4 Gb Myrinet (optical)
- 10/100/1000 Ethernet (copper)
- KVM (copper)
- Remote console manager (RJ-45/CAT5)
- Local console manager (CAT5E)
- Topspin 120 InfiniBand server (RS232)
- Topspin InfiniBand host channel adapter
- Voltaire 9024 InfiniBand switch router
- Copper Pass-Thru Module (copper)

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

Note: 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.

1 Gigabit Ethernet cabling

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

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.

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.

Fibre-channel cabling

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

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 CAT5 KVM cable from a cabinet containing cluster nodes to the cabinet containing the RCM, use a CCO cable and a CAT5 cable sufficiently long enough to reach the RCM switch.
- Multiple KVM switches can be connected in series.
- Up to 40 cluster nodes (xSeries 336 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 336 server) can be connected in series.

Topspin InfiniBand switch and host channel adapter cabling

Each Topspin 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.

Voltaire InfiniBand switch router cabling

Each Voltaire InfiniBand switch accepts one RS-232 cable for the Ethernet management port.

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 wire ties 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 4. Checking the Cluster 1350 cabling process

The IBM @server Cluster 1350 comes without an operating system installed unless the option is purchased that includes a preloaded operating system. An IBM customer service representative (CSR) performs the hardware cabling process then the customer performs the operating system and device 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.

Note: If you have a component that does not display the operational blue LEDs on the control panel, see “Diagnosing a lights out or brown out event” on page 32.

Installing cables and checking connections in the expansion cabinets

To install cables and check the connections in the expansion cabinets, complete the following steps:

1. Make sure that the circuit-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 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.

Installing cables and checking connections in the primary cabinet

To install cables and check the connections in the primary cabinet, complete the following steps:

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.
 - 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

To turn on the expansion cabinets, complete the following steps:

1. Turn on the cluster nodes using their power switches.
2. Make sure that every RSA 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

To turn on the primary cabinet, complete the following steps:

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 CAT5 and fibre channel connections have a green link LED.
 7. Make sure that each RSA 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 RSA 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 hardware installation

A startable CD is included with the Linux Cluster Installation Tool (LCIT) installation materials. With the 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 the cluster to make sure that the two sets of tab files are identical. If the tab files are not identical recheck the component installation. You can also use LCIT to help you locate RSAs or nodes that do not start as expected.

The LCIT interface has three information panels, each information panel is a tear-away panel that can be relocated or hidden.

- **Tree view** - The information in the left panel of the LCIT program window displays cluster device information and is presented in a collapsible tree view.
- **Information view** - The information in the right panel of the LCIT program window displays context-sensitive device information and changes depending on what you select in the tree view. For example, if you click a controller in the tree view, the Controller window opens in the information panel and allows you to review device information or to make changes on that controller.
- **Log view** - The log panel is located at the bottom of the LCIT program window and displays the most recent log messages.

To verify the installation of the cluster hardware, complete the following steps:

1. Make sure that the BIOS boot order is configured to boot the CD-ROM drive first.

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 3.3**. The LCIT window opens.
4. From the **View** menu, select **Discover**. The Cluster Management Node information displays in the information view.
5. Click the **Search Ranges** tab, and make sure that the IP address ranges are appropriate and if applicable, create a subset of the subnet.
6. Click the **IP Addresses** tab and click **Scan**. The discovery process begins and the tree view is dynamically populated with a list of the discovered cluster devices.
7. In the tree view, click each Ethernet controller and visually inspect the physical device to see if the nodes associated with that controller light their blue LEDs. When you click on a specific node, the blue LED on that node remains lit and all the other associated nodes blue LEDs turn off. You can use this to diagnose failing components.

Note: Devices that are listed in red have duplicate IP address, a duplicate name, or other issue that needs to be resolved.

8. In the tree view, click an Ethernet controller, and in the information view, click **Power On** or **Power Off** to recycle the power state of the controller. Repeat this step for each Ethernet controller in the cluster.

Note: If the current power state of the Ethernet controller is On, a red circle is displayed next to the controller in the tree view. If the current power state is Off, a black circle is displayed next to the controller. If the current power state is indeterminate, no circle is displayed next to the controller. The Ethernet controllers display dynamic statistical information.

Diagnosing a lights out or brown out event

If you experience a problem after the hardware installation, you can do a visual check of the cluster components to see if the cluster components light the various component LED indicators after installation. The following sequence occurs during a lights out or brown-out event scenario.

1. A lights-out event or brown-out event occurs. The system turns off and then turns back on through an external source.
2. 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.
3. Log files show system restart events on nodes and on RSAs. 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)
 - RSA event log
 - BIOS code event log

For more information about a lights out event, go to Chapter 6, “Troubleshooting hardware and software problems,” on page 37 or see the *IBM @server Cluster 1350 Linux Cluster Installation Tool User’s Guide* for more information about using LCIT to view the various log messages.

Related topics

See Appendix C, “Error and event logs,” on page 59.

Chapter 5. Accessing the cluster from a remote location

This chapter includes information about:

- accessing the cluster from a remote location
- accessing each node before the operating system is installed

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

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 unit 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 `rconsole` command accesses the remote console. 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]]  
[-f] [-o] | -O number of columns]] [-c [-v] | [-r | -f]]
```

Related topics

- Chapter 4, “Checking the Cluster 1350 cabling process,” on page 29
- Appendix C, “Error and event logs,” on page 59

Chapter 6. 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 such as server nodes and associated firmware, storage and networking subsystems, it also includes CSM software and General Parallel File System (GPFS) software.

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

This chapter includes information for the diagnosis of problems 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.

Displaying node configuration information

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 before installing LCIT software, at the console prompt, type:

```
Isnode -l
```

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 RSA 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 12.

Table 12. Troubleshooting the shared VLAN

Symptom	Action
<ul style="list-style-type: none"> • 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. 	<ol style="list-style-type: none"> 1. Make sure that the links between the management node, storage nodes, Cisco 3750, and 400x switches are turned on. 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, <code>ifconfig</code>. 4. Check the Cisco (3750 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 12 did not correct the problem, continue with the steps shown in “Clustering with one network.”

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.

To resolve the problem, complete the following steps:

1. Telnet or use (shell script) SSH to connect to the node through the serial console or KVM and make sure that 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 48 and complete the steps in that table to resolve the problem.
 - 2) If there are no errors, this indicates a network problem. See Table 13 on page 39 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 42 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 13 on page 39 and complete the listed actions to resolve the problem.
- Network adapter on the management node
- DHCP configuration
- Network configuration
- Cisco blade failure

Table 13. 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.	<ol style="list-style-type: none"> 1. At the console prompt, type the <code>ifconfig</code> command to make sure that the IP settings are correct. 2. Make sure that the cables are fully connected into the switch and node, and that everything else is connected 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.

To resolve the problem, complete the following steps:

1. 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.
2. If there are errors, go to “Isolating software problems” on page 48 and complete the listed actions to resolve the problem.
3. If there are no errors, this indicates a network problem. Go to Table 15 on page 41 and complete the listed actions to resolve the problem.
4. If Telnet or SSH fails, this indicates a node hardware problem. Go to “Isolating hardware problems” on page 42 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 14.

Table 14. Factory defaults

Device	IP address	Host name
Management node	<ul style="list-style-type: none"> • cluster VLAN: 172.20.0.1 • management VLAN: 172.30.0.1 	<ul style="list-style-type: none"> • eth0 – mgt.cluster.com • eth1– mgt1.cluster.com

Table 14. Factory defaults (continued)

Management node alias	172.29.101.1	bmc
e326 node bmcs	172.29.0.1	bmc001
Storage node	172.20.1.1	storage001
DS4000 storage controller	172.30.2.1	triton001
xSeries 336 cluster node	172.20.3.1	node001...nodexxx
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
RSA (bottom device)	172.30.30.1	rsa001
RSA (next device)	172.30.30.2	rsa002
RSA (Myrinet switch)	172.30.30.3	rsa003
Cisco Gigabit Ethernet switches	172.30.50.1	<ul style="list-style-type: none"> • cisco3508-001 • cisco3750-001
SMC switch	172.30.50.1	smc8624-001
Cisco 6500 series switch	172.30.80.1	<ul style="list-style-type: none"> • cisco6503-001 • cisco6509-001
Cisco 4006 switch (console management)	172.30.80.1	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 48 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 42 for problem resolution.

Table 15. Network troubleshooting for a cluster with two networks

Symptom	Action
Cannot ping a node or nodes on the cluster network from the management node, but the rconsole command and access from the KVM work correctly.	<ol style="list-style-type: none">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 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.

Isolating hardware problems

This section includes troubleshooting information about cluster hardware problems.

Node checks

Table 16. Troubleshooting the remote console network

Symptom	Action
<ol style="list-style-type: none">1. Cannot execute a rconsole command to any cluster node.2. Cannot execute any rconsole commands to get an active terminal session.	<ol style="list-style-type: none">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. Replace and configure the terminal server, 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 that the serial port (COM 1) is configured to redirect output to the terminal server, complete the following steps:<ol style="list-style-type: none">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:<ul style="list-style-type: none">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 336 and xSeries 346 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 16 on page 42 do not correct the issue you may have a problem with a port on the terminal server.

To test a different port, complete the following steps:

1. At the console prompt, type the following command and record the port information:

```
lsmode -I <nodename> lgrep Port
```

2. Move the cable to a new port and change the port number; then, type
chnode <nodename> ConsolePortNum=xx

where *xx* is the new port number.

If the symptom persists, go to “Checking service processor logs” on page 50 and check the service processor log.

Hardware problem in service processor log

Go to “Node checks” on page 42 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 RSA. Otherwise, go to “Node checks” on page 42 for node problem resolution.

rpower to node fails

To resolve remote power failures, use the following information:

- Check the service processor connection.
- At the console prompt, type **rpower -a on**.
- Go to “Checking service processor logs” on page 50 and check the service processor log.
- Use the Web interface, Telnet, or SSH to reach each RSA adapter on the cluster and then connect to the service processor on each cluster node individually through the Remote Advanced System Management (ASM) Access menu.
- If the cluster node is not 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 17. Troubleshooting the service processor network

Symptom	Action
<ol style="list-style-type: none"> 1. The <code>rpower -a</code> query command does not return with the status of all nodes 2. Cannot see all the nodes when managing remote ASM service processors. 3. Cannot connect to individual RSA adapters using a Web browser. 	<ol style="list-style-type: none"> 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 adapter. 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 CAT5 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 336 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 3750 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 17 did not correct the problem, continue with the steps shown in "Remote Supervisor Adapter connection failure."

Remote Supervisor Adapter connection failure

To resolve RSA adapter connection failures, complete the following steps:

1. To make sure that the node has power, type:

```
rpower query
```

- a. If the node has power, ping the RSA using the `HWControlPoint` field in the `l snode` output.

- 1) If ping succeeds, reset the RSA. To reset the RSA, 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 at the Ethernet switch, go to "Resetting the Remote Supervisor Adapter" on page 50 and reset the RSA adapter.

Node connection or command failure

If the RSA connection is working, but the node connection or commands issued to the node failed:

- Connect to the RSA adapter and check the node list.
- Check all cabling.
- Go to “Node checks” on page 42 and perform node checks.

Checking storage

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

Table 18. Troubleshooting the fibre-storage network

Symptom	Action
Cannot see disk drives from the storage node.	<ol style="list-style-type: none">1. Reboot the storage node and press the Alt + Q keys to go into QLogic setup. Make sure that the DS4 500 is a listed device.2. Check the fibre connections between the server host bus adapter (HBA) and the hubs on the DS4 500. The green connection LED should be lit.3. Check cabling on the DS4 500 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 per second.4. Check the blade server and Enhanced System Manager (ESM) firmware levels and if necessary, update them to the current supported levels.

If following the steps in Table 18 did not correct the problem, continue with the steps shown in “File system failure.”

File system failure

Check the disks using fdisk -l:

- If fdisk -l completes without error, go to “GPFS checks” on page 49 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 complete 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 documentation that comes with the switch and complete the problem resolution process.

- If the node does not respond, go to “Node checks” on page 42 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 49 and continue with Myrinet problem resolution.
- If the node does not respond, go to “Node checks” on page 42 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 42 and continue with Node checks.
- If following the steps in Table 19 do not correct the problem, continue with the steps shown in “Isolating network, node, and Linux problems” on page 37.

Table 19. Troubleshooting the terminal server network for the Remote Console

Symptom	Action
Unable to run rconsole commands to get an active terminal session.	<ol style="list-style-type: none"> 1. Make sure that all cables and connectors are seated correctly at the nodes and the In-Reach terminal server. 2. Ping the unit at 172.30.20.1. 3. Replace and configure the new terminal server, then at the *IN-Reach_Priv>* prompt, type, <i>show port <portnumber></i>. 4. Ping 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: <ul style="list-style-type: none"> • Restart the node and 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. • Configure 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 20. Troubleshooting the KVM network

Symptom	Action
The KVM switch selector shows some or all systems are non-active (indicated by a red X) but the system is turned on.	<ol style="list-style-type: none">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 CAT5 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 CAT5 (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 49 and continue with the file system problem resolution.
- If **fdisk -l** reports missing disks, make sure that the adapter device driver is configured:
 - If the adapter device driver is configured, go to “Checking storage” on page 45 and continue with storage subsystem problem resolution.
 - If the adapter device driver is not configured, check the adapter hardware and then go to the Linux documentation that comes with your software and resolve the configuration issue.

PFA alert indicates internal disk

Go to “Checking storage” on page 45 and complete 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 Linux documentation that comes with the 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 Linux documentation that comes with the 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 documentation that comes with CSM to 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 documentation that comes with CSM to 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 documentation that comes with CSM to complete the problem resolution process.

RMC not running

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

Commands or file replication fails

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

rpower or rconsole commands fail

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

GPFS checks

Performance problems

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

GPFS file system failure

See the *GPFS problem resolution* and *GPFS Performance White papers* included with the 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 `lnode -l` command to determine the host name for the RSA adapter and the service processor name associated with the failing node. Use the Telnet command, SSH command, or a Web browser to connect to the RSA 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

The RSA adapter is typically connected to a remote power control strip.

To reset the RSA, complete the following steps:

Plug the remote power control strip of the failing RSA into another power source and type the power off and power on commands to the RSA port.

OR

1. Start LCIT. LCIT opens.
2. From the **View** menu, select **Discover**.
3. In the tree view, click an Ethernet Controller expand icon to list the associated RSA nodes. A list of devices opens.
4. Click the RSA node that you want to reset. The Node window opens.
5. Click the **System** tab and complete one of the following actions:
 - To turn on the power to all devices in the RSA chain, click **Power On**.
 - To turn off the power to all devices in the RSA chain, click **Power Off**.
 - To turn the power on and off to all devices in the RSA chain, click **Reset**.
 - To reset all ISMP processors in the RSA chain, click **Reset Remote ISMPs**. This turns off the power for all ISMP devices in the RSA chain then turns the power on.
 - To reset and reboot an RSA device, click **Reset RSA Device**.

Checking service processor logs

To check service processor logs, complete one of the following procedures:

At the console prompt, type, **l snode -1** to determine the host name for the RSA 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 using the host name, and select **View Log**.

OR

1. Start LCIT. LCIT opens.
2. From the **View** menu, select **Discover**.
3. In the tree view, click an Ethernet Controller expand icon to list the associated RSA nodes. A list of devices opens.

4. Click the RSA node whose logs you want to view. The Node window opens.
5. Click the **Logs** tab.

Management, cluster, and storage node problems

The IBM components used for management, cluster, and storage nodes are shown in Table 21:

Table 21. IBM components used for management, cluster, and storage nodes

Node type	IBM component used
Management node	<ul style="list-style-type: none"> • xSeries 346 • @server 326 (64-bit operating system environment)
Cluster node	<ul style="list-style-type: none"> • @server 326 • xSeries 336 • xSeries 346 • BladeCenter
Storage node	<ul style="list-style-type: none"> • @server 326 • xSeries 336 (do not use in a Fibre-Channel storage configuration) • xSeries 346 • xSeries 360

While the @server 326, xSeries 336, xSeries 346, and xSeries 360 are all high-reliability units, occasionally a component may fail. Two areas that might cause problems are:

- Disk drives
- System board

The following section includes information about:

- Disk drive failures
- System board failures
- xSeries 336 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 346 supports hot swapping of hard disks. To replace a failing hard disk on the management node, complete the following steps:

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 the *IBM @server Linux Cluster Installation Tool User's Guide* 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 336 supports hot swapping of hard disks, but the BladeCenter unit does not. To replace a failing hard disk on an xSeries 336:

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, type:

```
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 unit, first power down the blade server. Next, remove the blade server from the BladeCenter unit and replace the hard drive as outlined in *IBM @server BladeCenter Installation and User's Guide for your Blade Server Type*.

Once the drive is replaced and the blade server is returned to the BladeCenter unit at the management node, type:

```
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 346 server and xSeries 360 server support hot swapping of hard disks.

To replace a failing hard disk on the storage node, complete the following steps:

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

To troubleshoot information about system board failures, complete the following steps:

1. Flash the system BIOS code to the level used in the installation. Refer to the *IBM @server Linux Cluster Installation Tool User's Guide* for a listing of the software and firmware levels used in the Cluster 1350.
2. Flash the diagnostics to match the BIOS code level. Refer to the *IBM @server Linux Cluster Installation Tool User's Guide* for a listing of the software and firmware levels used in the Cluster 1350.
3. Flash the onboard ASM to the current level. Refer to the *IBM @server Linux Cluster Installation Tool User's Guide* for a listing of the software and firmware levels used in the Cluster 1350.
4. Configure the following:
 - 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 RSA is installed in this node (xSeries 336, xSeries 346, xSeries 360 only).
- If you are replacing an HS20 BladeCenter unit 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 server card you installed.
 - To get the MAC address of eth0 for the new component, type:


```
getadapters -wn <nodename>
```

Contact support for any setup or IP configurations that need to be performed.

xSeries 336 problems

In a xSeries 336 with an RSA over C2T connection make sure that the cluster node at the beginning of the C2T chain has an RSA, 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 unit 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 connected properly. An improper switchblock setting, loose option card, or disconnected 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 unit then you must reassign the IP address for the external ports to work. Make sure that 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 a 220V power source. BladeCenter units connected to PDUs connected to 115V power will not function correctly.

Power problems

The following section includes information about:

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

No power to multiple devices

To troubleshoot power failures to multiple devices, complete the following steps:

1. Make sure that the 30 amp twist-lock plugs are locked into the customer supplied receptacles.
2. Make sure that the main power circuit breakers at the customer circuit breaker panel are in the ON position.
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 facility 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 PDU breakers are in the ON position.
5. Make sure that the PDU plugs are securely seated into the Power Out sockets on the Frame Power Blocks.
6. Make sure that there is voltage at the Power Out ports on the PDU by 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 there is voltage at the Power Out ports on the PDU by 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 light up on the failing unit, replace the power supply or complete unit if the power supply cannot be replaced.

Related documentation

Additional hardware maintenance and problem resolution information relating to the xSeries 336 and xSeries 346 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. To check for updated documentation and technical updates, complete the following steps:

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.

Appendix A. Getting help and technical assistance

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. 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 all cables to make sure that they are connected.
- Check the power switches to make sure that the system is turned on.
- 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/>.
- 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 17X rack model and serial number is generally located in the rear, on the bottom-right.
 - The 1417 model 11X rack model and serial number is generally located in the rear, on the bottom-right.

Using the documentation

Information about your IBM system and preinstalled software, if any, is available in the documentation that comes with your system. That documentation includes printed books, online books, readme files, and help files. See the troubleshooting information in your system documentation for instructions for using the diagnostic programs. The troubleshooting information or the diagnostic programs might tell you

that you need additional or updated device drivers or other software. IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. To access these pages, go to <http://www.ibm.com/pc/support/> and follow the instructions. Also, you can order publications through the IBM Publications Ordering System at <http://www.elink.ibm.com/public/applications/publications/cgibin/pbi.cgi>.

Getting help and information from the World Wide Web

On the World Wide Web, the IBM Web site has up-to-date information about IBM products, services, and support. The address for IBM xSeries information is <http://www.ibm.com/eserver/xseries/>.

You can find service information for your IBM products, including supported options, at <http://www.ibm.com/pc/support/>.

Software service and support

Through IBM Support Line, you can get telephone assistance, for a fee, with usage, configuration, and software problems with xSeries servers, IntelliStation[®] workstations, and appliances. For information about which products are supported by Support Line in your country or region, go to <http://www.ibm.com/services/sl/products/>.

For more information about Support Line and other IBM services, go to <http://www.ibm.com/services/>, or go to <http://www.ibm.com/planetwide/> for support telephone numbers. In the U.S. and Canada, call 1-800-IBM-SERV (1-800-426-7378).

Hardware service and support

You can receive hardware service through IBM 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.

Appendix B. Frequently asked questions

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

Question: Why do I sometimes get the error message *2651-689 Java interface error for method query: SPEXception?*

Answer: This is due to a defect in the 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`

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

Answer: Sometimes the RSA adapters get hung. They can be reset using the Web interface, Telnet to the RSA adapter, or issuing the command `rpower -n <node> resetsp_hcp`

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

Answer: tftp is not loaded or configured on the management node.

Question: Why doesn't the xSeries 346 boot PXE correctly?

Answer: You cannot have a PCI ethernet card that uses the e1000 driver in the xSeries 346 when installing. Take the card out and retry the installation.

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

Answer: 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.

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

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

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

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

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

Answer: 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 *csmssetupast* command 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. Using the console, watch the boot screen and the error messages should be on the console when the process has completed.

Question: Why do SLES installs take forever?

Answer: 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.

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

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

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 `eventlog -n <nodename> all` command.

Remote Supervisor Adapter log

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.`

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 */ftpboot* 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 Controller instead of the local SCSI drive connected to the internal Adaptec SCSI controller. Why this happens is as follows:

When the modules are loaded, the order ends up in such a way that the driver for the QLogic fibre-channel 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 that the KVM switch that was added is in default settings mode.

RSA and service processor

If there are any RSA errors, check to make sure that 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 or 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, reseal 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 `lsmmod` command.
- Check to see if the Myricom adapter is up and running by using the `ifconfig` command.

Appendix E. 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 RSA 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 than 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.

For more information on switch commands, see the documentation that comes with the device.

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 verify the operation of any non-IBM product, program, or service.

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

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

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 2005. 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	Predictive Failure Analysis
Active PCI	PS/2
Active PCI-X	ServeRAID
Alert on LAN	ServerGuide
BladeCenter	ServerProven
C2T Interconnect	TechConnect
Chipkill	ThinkPad
EtherJet	Tivoli
e-business logo	Tivoli Enterprise
@server	Update Connector
FlashCopy	Wake on LAN
IBM	XA-32
IBM (logo)	XA-64
IntelliStation	X-Architecture
NetBAY	XcelL4
Netfinity	XpandOnDemand
NetView	xSeries
OS/2 WARP	

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 of Sun Microsystems, Inc. in the United States, other countries, or both.

Adaptec and HostRAID are trademarks of Adaptec, Inc., in the United States, other countries, or both.

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

Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

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>.

IBM encourages owners of information technology (IT) equipment to responsibly recycle their equipment when it is no longer needed. IBM offers a variety of programs and services to assist equipment owners in recycling their IT products. Information on product recycling offerings can be found on IBM's Internet site at <http://www.ibm.com/ibm/environment/products/prp.shtml>.



Notice: This mark applies only to countries within the European Union (EU) and Norway.

This appliance is labeled in accordance with European Directive 2002/96/EC concerning waste electrical and electronic equipment (WEEE). The Directive determines the framework for the return and recycling of used appliances as applicable throughout the European Union. This label is applied to various products to indicate that the product is not to be thrown away, but rather reclaimed upon end of life per this Directive.

注意：このマークは EU 諸国およびノルウェーにおいてのみ適用されます。

この機器には、EU 諸国に対する廃電気電子機器指令 2002/96/EC(WEEE) のラベルが貼られています。この指令は、EU 諸国に適用する使用済み機器の回収とリサイクルの骨子を定めています。このラベルは、使用済みになった時に指令に従って適正な処理をする必要があることを知らせるために種々の製品に貼られています。

Remarque : Cette marque s'applique uniquement aux pays de l'Union Européenne et à la Norvège.

L'étiquette du système respecte la Directive européenne 2002/96/EC en matière de Déchets des Equipements Electriques et Electroniques (DEEE), qui détermine les dispositions de retour et de recyclage applicables aux systèmes utilisés à travers l'Union européenne. Conformément à la directive, ladite étiquette précise que le produit sur lequel elle est apposée ne doit pas être jeté mais être récupéré en fin de vie.

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 return 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.



For Taiwan: Please recycle batteries.



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.

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

聲 明
此為 A 級產品。在生活環境中，該產品可能會造成無線電干擾。在這種情況下，可能需要用戶對其干擾採取切實可行的措施。

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。

Power cords

For your safety, IBM provides a power cord with a grounded attachment plug to use with this IBM product. To avoid electrical shock, always use the power cord and plug with a properly grounded outlet.

IBM power cords used in the United States and Canada are listed by Underwriter's Laboratories (UL) and certified by the Canadian Standards Association (CSA).

For units intended to be operated at 115 volts: Use a UL-listed and CSA-certified cord set consisting of a minimum 18 AWG, Type SVT or SJT, three-conductor cord, a maximum of 15 feet in length and a parallel blade, grounding-type attachment plug rated 15 amperes, 125 volts.

For units intended to be operated at 230 volts (U.S. use): Use a UL-listed and CSA-certified cord set consisting of a minimum 18 AWG, Type SVT or SJT, three-conductor cord, a maximum of 15 feet in length and a tandem blade, grounding-type attachment plug rated 15 amperes, 250 volts.

For units intended to be operated at 230 volts (outside the U.S.): Use a cord set with a grounding-type attachment plug. The cord set should have the appropriate safety approvals for the country in which the equipment will be installed.

IBM power cords for a specific country or region are usually available only in that country or region.

IBM power cord part number	Used in these countries and regions
02K0546	China
13F9940	Australia, Fiji, Kiribati, Nauru, New Zealand, Papua New Guinea
13F9979	Afghanistan, Albania, Algeria, Andorra, Angola, Armenia, Austria, Azerbaijan, Belarus, Belgium, Benin, Bosnia and Herzegovina, Bulgaria, Burkina Faso, Burundi, Cambodia, Cameroon, Cape Verde, Central African Republic, Chad, Comoros, Congo (Democratic Republic of), Congo (Republic of), Cote D'Ivoire (Ivory Coast), Croatia (Republic of), Czech Republic, Dahomey, Djibouti, Egypt, Equatorial Guinea, Eritrea, Estonia, Ethiopia, Finland, France, French Guyana, French Polynesia, Germany, Greece, Guadeloupe, Guinea, Guinea Bissau, Hungary, Iceland, Indonesia, Iran, Kazakhstan, Kyrgyzstan, Laos (People's Democratic Republic of), Latvia, Lebanon, Lithuania, Luxembourg, Macedonia (former Yugoslav Republic of), Madagascar, Mali, Martinique, Mauritania, Mauritius, Mayotte, Moldova (Republic of), Monaco, Mongolia, Morocco, Mozambique, Netherlands, New Caledonia, Niger, Norway, Poland, Portugal, Reunion, Romania, Russian Federation, Rwanda, Sao Tome and Principe, Saudi Arabia, Senegal, Serbia, Slovakia, Slovenia (Republic of), Somalia, Spain, Suriname, Sweden, Syrian Arab Republic, Tajikistan, Tahiti, Togo, Tunisia, Turkey, Turkmenistan, Ukraine, Upper Volta, Uzbekistan, Vanuatu, Vietnam, Wallis and Futuna, Yugoslavia (Federal Republic of), Zaire
13F9997	Denmark
14F0015	Bangladesh, Lesotho, Macao, Maldives, Namibia, Nepal, Pakistan, Samoa, South Africa, Sri Lanka, Swaziland, Uganda

IBM power cord part number	Used in these countries and regions
14F0033	Abu Dhabi, Bahrain, Botswana, Brunei Darussalam, Channel Islands, China (Hong Kong S.A.R.), Cyprus, Dominica, Gambia, Ghana, Grenada, Iraq, Ireland, Jordan, Kenya, Kuwait, Liberia, Malawi, Malaysia, Malta, Myanmar (Burma), Nigeria, Oman, Polynesia, Qatar, Saint Kitts and Nevis, Saint Lucia, Saint Vincent and the Grenadines, Seychelles, Sierra Leone, Singapore, Sudan, Tanzania (United Republic of), Trinidad and Tobago, United Arab Emirates (Dubai), United Kingdom, Yemen, Zambia, Zimbabwe
14F0051	Liechtenstein, Switzerland
14F0069	Chile, Italy, Libyan Arab Jamahiriya
14F0087	Israel
1838574	Antigua and Barbuda, Aruba, Bahamas, Barbados, Belize, Bermuda, Bolivia, Brazil, Caicos Islands, Canada, Cayman Islands, Costa Rica, Colombia, Cuba, Dominican Republic, Ecuador, El Salvador, Guam, Guatemala, Haiti, Honduras, Jamaica, Japan, Mexico, Micronesia (Federal States of), Netherlands Antilles, Nicaragua, Panama, Peru, Philippines, Taiwan, United States of America, Venezuela
24P6858	Korea (Democratic People's Republic of), Korea (Republic of)
34G0232	Japan
36L8880	Argentina, Paraguay, Uruguay
49P2078	India
49P2110	Brazil
6952300	Antigua and Barbuda, Aruba, Bahamas, Barbados, Belize, Bermuda, Bolivia, Caicos Islands, Canada, Cayman Islands, Colombia, Costa Rica, Cuba, Dominican Republic, Ecuador, El Salvador, Guam, Guatemala, Haiti, Honduras, Jamaica, Mexico, Micronesia (Federal States of), Netherlands Antilles, Nicaragua, Panama, Peru, Philippines, Saudi Arabia, Thailand, Taiwan, United States of America, Venezuela

Appendix G. Declaration of IBM product noise emission values

General noise emission values for the Cluster 1350 are shown in Table 22.

Table 22. General noise emission values for the Cluster 1350

Declaration of IBM product noise emission values				
Product configuration	Declared A-weighted sound power levels L_{wAd} (B)		Declared A-weighted sound pressure level L_{pAm} (db) at 1 meter bystander position	
Cluster 1350 (fully configured racks)	8.0 (bels) operating	8.0 (bels) idling	62 (db) operating	62 (db) idling
Note: L_{wAd} is the declared (upper limit) sound power level for a random sample of machines. L_{pAm} is the mean value of the A-weighted sound pressure levels at the bystander (1 - meter) positions for a random sample of machines. All measurements are made in accordance with ISO 7779 and declared in conformance with ISO 9296.				

Noise emission values are different for racks that contain BladeCenters. Each BladeCenter unit has two blowers for cooling. An optional acoustics module is available that dampens the noise generated by the blowers. Table 23 shows the maximum sound levels emitted by BladeCenter units without the acoustics module and with the acoustics module. This represents the minimum and maximum number of BladeCenter units that can be installed in a 42-U rack.

Table 23. Noise emission values for racks containing BladeCenters

	Number of BladeCenter units per rack	Maximum sound power idle	Maximum sound power operating
Without acoustics module	1	7.4 bels	7.4 bels
With acoustics module	1	6.9 bels	6.9 bels
Without acoustics module	4	8 bels	8 bels
With acoustics module	4	7.5 bels	7.5 bels
Without acoustics module	6	8.2 bels	8.2 bels
With acoustics module	6	7.7 bels	7.7 bels

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.

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

Numerics

- 1 Gigabit Ethernet cabling 25
- 1 U flat-panel console
 - cluster components 7
- 10/100/1000 Ethernet switch
 - cluster components 8
 - description 8
- 2x8 console switch
 - KVM switch description 7

A

- access
 - remote 35
 - remote console 35
 - remote power command 35
- APC event log
 - viewing 59

B

- BladeCenter Ethernet switch module
 - factory default address 40
- BladeCenter problems
 - problem determination 53
- blue CAT5E Ethernet cables
 - part number 13

C

- cabinet connections
 - checking 29
- cabinet placement
 - service clearances 11
- cable
 - replacing defective harness 26
- cables
 - color-coded Ethernet cables 13
- cabling 17
 - 1 Gigabit Ethernet 25
 - 10/100/1000 Ethernet switch 25
 - connectors 29
 - Fibre Channel 25
 - high-speed 10/100/1000 Ethernet switch 25
 - intercabinet, general information 17
 - intracabinet 17
 - intracabinet, general information 17
 - KVM switch 25
 - local console manager 26
 - Myrinet switch 25
 - overview 17
 - RCM 26
 - Remote Console Manager 26
 - Topspin switch 26
 - types of intercabinet 12, 24
 - Voltaire switch 26
- cabling, intercabinet 17

- CAT5E cables
 - color-coded Ethernet cables 13
 - part numbers 13
- checking connections
 - primary cabinet 30
- checking expansion unit
 - connectors 29
- checking primary unit
 - connectors 30
- Cisco 4000 series switch
 - factory default address 40
- Cisco 6500 series switch
 - factory default address 40
- Cisco Gigabit Ethernet switches
 - factory default address 40
- Cisco line card
 - description 9
- Class A electronic emission notice 71
- clearances
 - floor loading 15
- cluster
 - lights out or brownout 32
 - power down 35
 - power up 29
 - turn on procedure 29
- cluster components
 - 1 U flat-panel console description 7
 - 10/100/1000 Ethernet switch 8
 - cluster node 5
 - expansion modules 8
 - Fibre Channel adapters and switches 6
 - high-speed Myrinet switch 8
 - InfiniBand switches and devices 7
 - KVM switch description 7
 - management node 5
 - Myrinet BladeCenter connectivity 9
 - power distribution unit 9
 - SCSI/RAID storage controller adapter 7
 - storage expansion unit 6
 - storage node 5
 - storage server 6
 - terminal server 8
 - Topspin Infiniband host channel adapters 7
- cluster management
 - administer 35
- cluster node
 - 326 4
 - 336 4
 - 346 4
 - Blade servers 4
 - cluster components 5
 - disk drive failure 52
- cluster node disk drive failure
 - troubleshooting 52
- color-coded Ethernet cables
 - CAT5E cables 13
 - part numbers 13

- console
 - description 7
- copper pass-thru module
 - description 8
 - expansion module 8
- CSM
 - display node configuration 37
 - known problems 61
 - problem determination 48
- CSM event log
 - viewing 59
- CSM node configuration
 - display 37

D

- defective cable harness
 - replacing 26
- description
 - 10/100/1000 Ethernet switch 8
 - Cisco line card 9
 - cluster nodes 5
 - console 7
 - copper pass-thru module 8
 - FASTt 250 GB dual-port serial ATA disk drive 6
 - FASTt EXP100 6
 - FASTt EXP400 6
 - FASTt U320 146 GB hard disk drive 6
 - FASTt U320 300 GB hard disk drive 6
 - FASTt600 storage controller 6
 - FASTt600 with Turbo storage controller 6
 - FASTt900 storage controller 6
 - high-speed Cisco switches 9
 - high-speed Myrinet switch 8
 - high-speed Myrinet switches 8
 - IBM Distributed Power Interconnect front-end power distribution unit 9
 - IBM Distributed Power Interconnect high-density power distribution unit 9
 - IBM Distributed Power Interconnect rack power distribution unit 9
 - IBM TotalStorage DS300 SAN storage server 6
 - IBM TotalStorage DS400 SAN storage server 6
 - information view 31
 - KVM switch 7
 - log view 31
 - management nodes 5
 - management VLAN 17
 - Maxtor 250 GB S-ATA hard disk drive 6
 - Myrinet BladeCenter connectivity 9
 - Myrinet M3-E128 model 8
 - Myrinet M3-E32 model 8
 - Myrinet M3-E64 model 8
 - Myrinet M3-SW16-8E model 8
 - Myrinet M3F2-PCIXE-2 model 8
 - Myrinet M3F2-PCIXE-4 model 8
 - NetBAY front-end power distribution unit 9
 - NetBAY rack-power distribution unit 9
 - optional secondary cluster VLAN 17
 - power distribution unit 9
 - primary cluster VLAN 17

- description (*continued*)
 - PRO/1000 MT dual port server adapter 8
 - Savvio 2.5 inch small form factor hard disk drive 6
 - serial ATA hard disk drive 6
 - ServeRAID -6I Ultra320 SCSI controller 7
 - ServeRAID -6M Ultra320 SCSI controller 7
 - ServeRAID 7k serial-ATA controller 7
 - ServeRAID 7t serial-ATA controller 7
 - SMC mini-GBIC expansion module 8
 - storage nodes 5
 - storage servers 6
 - terminal server 8
 - Topspin 12-port 4X InfiniBand line card 7
 - Topspin 120 InfiniBand switch 7
 - Topspin 270 InfiniBand management module 7
 - Topspin 270 InfiniBand switch 7
 - Topspin 270 InfiniBand switch fabric module 7
 - Topspin InfiniBand host channel adapter 7
 - Topspin InfiniBand host channel adapter
 - PCI-Express 7
 - Topspin InfiniBand host channel adapter PCI-X 7
 - tree view 31
 - Voltaire 9024 InfiniBand switch 7
- determining problems 37
- disk drive failure
 - cluster node 52
 - troubleshooting, storage node 52
- display
 - CSM node configuration file 37
- DPI
 - description 9

E

- e326 node bmc
 - factory default address 40
- electronic emission Class A notice 71
- emission values, product noise 75
- engaging
 - outriggers 15
- error logs 59
- eserver 326
 - management node 5
- Ethernet cabling, 1 Gigabit 25
- event
 - logs 59
- expansion cabinets
 - turning on 30
- expansion module
 - copper pass-thru module 8
 - PRO/1000 MT dual port server adapter 8
 - SMC mini-GBIC expansion module 8
- expansion modules
 - cluster components 8

F

- factory default address
 - BladeCenter Ethernet switch module 40
 - Cisco 4000 series switch 40
 - Cisco 6500 series switch 40

- factory default address *(continued)*
 - Cisco Gigabit Ethernet switches 40
 - e326 node bmc 40
 - FASTt storage controller 40
 - management node 39
 - management node alias 40
 - Myrinet switch 40
 - remote console manager 40
 - Remote Supervisor Adapter 40
 - SMC switch 40
 - storage node 40
 - terminal server 40
 - Topspin 120 InfiniBand switch 40
 - xSeries 336 cluster node 40
- FAQ 57
- FASTt 250 GB dual-port serial ATA disk drive
 - description 6
- FASTt EXP100
 - storage expansion unit 6
- FASTt EXP400 storage expansion unit
 - description 6
- FASTt storage controller
 - factory default address 40
- FASTt U320 146 GB hard disk drive
 - description 6
- FASTt U320 300 GB hard disk drive
 - description 6
- FASTt600 storage controller
 - description 6
- FASTt600 with Turbo storage controller
 - description 6
- FASTt900 storage controller
 - description 6
- FCC Class A notice 71
- Fibre Channel adapters and switches
 - cluster components 6
- Fibre Channel cabling 25
- finding documentation 10
- floor loading
 - clearances 15
- Force 10 switch fabric module
 - supported Fibre Channel adapters and switches 6
- frequently asked questions 57

G

- GPFS
 - problem determination 49
- green CAT5E Ethernet cables
 - part number 13
- guidelines
 - placing the cabinets 13

H

- handling static-sensitive devices
 - directions xv
- hardware
 - problem determination 37
- hardware/software
 - problem determination 37

- high-speed 10/100/1000 Ethernet switch
 - cabling 25
- high-speed Cisco switches
 - description 9
- high-speed Myrinet switch
 - cluster components 8

I

- IBM Distributed Power Interconnect front-end power
 - distribution unit
 - description 9
- IBM Distributed Power Interconnect high-density power
 - distribution unit
 - description 9
- IBM Distributed Power Interconnect rack power
 - distribution unit
 - description 9
- IBM TotalStorage DS300 SAN storage server
 - description 6
- IBM TotalStorage DS400 SAN storage server
 - description 6
- IBM xSeries 336 and xSeries 346 51
- InfiniBand switches and line cards
 - cluster components 7
- information
 - intercabinet cabling 17
 - intracabinet cabling 17
- information view
 - description 31
- installation
 - rack 13
- installer responsibilities
 - installing the rack 13
- installing
 - rack stability kit 15
 - stabilizer kit 15
- intercabinet cabling 17
 - general information 17
 - types 12, 24
- interface
 - overview 31
- intracabinet cabling 17
 - general information 17

K

- known problems 61
 - CSM 61
 - KVM 62
 - Myrinet 63
 - node 61
 - RSA 63
 - service processor 63
 - storage 61
- KVM
 - description 7
 - known problems 62
- KVM switch
 - cabling 25
 - cluster components 7

- KVM switch description
 - 2x8 console switch 7
 - local console switch 7
 - remote console switch 7

L

- LCIT
 - verifying the installation 31
- LCIT event log
 - viewing 59
- LCIT installation
 - verifying 31
- Linux event log
 - viewing 59
- local console manager
 - cabling 26
- local console switch
 - KVM switch description 7
- locating
 - machine serial number 55
 - model number 55
 - rack serial number 55
 - serial number 55
- log view
 - description 31
- logs
 - error 59
 - event 59

M

- machine model number
 - locating 55
- machine serial number
 - locating 55
- management
 - cluster 35
- management node
 - components 5
 - eserver 326 5
 - factory default address 39
 - xSeries 346 5
- management node alias
 - factory default address 40
- management VLAN
 - description 17
- Maxtor 250 GB S-ATA hard disk drive
 - description 6
- model number
 - locating 55
- Myrinet
 - known problems 63
- Myrinet BladeCenter connectivity
 - cluster components 9
 - description 9
- Myrinet CLOS 256 switch
 - placement of rack 12
- Myrinet high-speed switch cabling 25
- Myrinet M3-E128 model
 - description 8

- Myrinet M3-E32 model
 - description 8
- Myrinet M3-E64 model
 - description 8
- Myrinet M3-SW16-8E model
 - description 8
- Myrinet M3F-PCIXD-2 model
 - description 8
- Myrinet M3F2-PCIXE-2 model
 - description 8
- Myrinet M3F2-PCIXE-4 model
 - description 8
- Myrinet switch
 - cabling 25
 - factory default address 40

N

- NetBAY front-end power distribution unit
 - description 9
- NetBAY rack-power distribution unit
 - description 9
- no power
 - individual device 54
 - no power to individual devices
 - power problems 54
 - no power to multiple devices
 - power problems 53
- node
 - known problems 61
- node event log
 - viewing 59
- nodes, xSeries 336 and xSeries 346 51
- notes, important 68
- notices
 - electronic emission 71
 - FCC, Class A 71
 - used in this book xv

O

- online documentation
 - locating 10
- online publications
 - locating 10, 54
- operating system support 1
- optional secondary cluster VLAN
 - description 17
- overview
 - cabling 17
 - interface 31
 - system 1
- overview, system 1

P

- part number
 - blue CAT5E Ethernet cables 13
 - green CAT5E Ethernet cables 13
 - grey CAT5E Ethernet cables 13
 - yellow CAT5E Ethernet cables 13

- placement of rack
 - Myrinet CLOS 256 switch 12
- placing the cabinets
 - guidelines 13
- power cords 73
- power distribution unit
 - cluster components 9
 - description 9
- power down
 - cluster 35
- power problems 53
 - no power to individual devices 54
 - no power to multiple devices 53
 - problem determination 53
- power up
 - cluster 29
- preparing
 - rear door heat eXchanger 16
- primary cabinet connections
 - checking 30
- primary cabinets
 - switching on 30
- primary cluster VLAN
 - description 17
- Pro/1000 MT dual port server adapter
 - description 8
- PRO/1000 MT dual port server adapter
 - expansion module 8
- problem determination 37
 - BladeCenter problems 53
 - checking service processor logs 50
 - cluster with one network 38
 - cluster with two networks 39
 - hardware 37
 - isolating hardware problems 42
 - isolating network, node, and Linux problems 37
 - isolating software problems 48
 - CSM 48
 - device driver not configured 48
 - GPFS 49
 - node unresponsive 48
 - power problems 53
 - resetting RSA cards 50
 - setting up SNMP alerts 49
 - SNMP monitoring 49
 - software 37
 - xSeries 336 53
 - xSeries 336 nodes 51
- problems
 - power 53
- procedure
 - lights out or brownout 32
- product noise emission values 75

R

- rack serial number
 - locating 55
- rack stability
 - engaging the outriggers 15

- rack support
 - engaging the outriggers 15
- RCM
 - cabling 26
- rear door heat eXchanger
 - preparing 16
- related publications
 - locating 10
- remote access
 - link 35
 - remote console 35
 - remote power command 35
- remote console manager
 - factory default address 40
 - KVM switch description 7
- Remote Console Manager
 - cabling 26
- Remote Supervisor Adapter
 - factory default address 40
- replacing
 - defective cable harness 26
- RSA
 - known problems 63
- RSA event log
 - viewing 59

S

- Savvio 2.5 inch small form factor hard disk drive
 - description 6
- SCSI/RAID storage controller adapter
 - cluster components 7
- serial number
 - locating 55
- ServeRAID 7k serial-ATA controller
 - description 7
- ServeRAID 7t serial-ATA controller
 - description 7
- service clearances
 - cabinet placement 11
- service processor
 - known problems 63
- SMC mini-GBIC expansion module
 - description 8
 - expansion module 8
- SMC switch
 - factory default address 40
- software
 - CSM
 - known problems 61
 - CSM problem determination 48
 - GPFS problem determination 49
 - problem determination 37
- stability kit
 - installation 15
- static-sensitive devices
 - directions for handling xv
- storage
 - known problems 61
- storage controller
 - FASTt600 description 6

- storage controller *(continued)*
 - FAST600 with Turbo description 6
 - FAST900 6
- storage controller adapter
 - description 7
- storage expansion unit
 - cluster components 6
 - FAST EXP100 description 6
 - FAST EXP400 description 6
- storage node
 - cluster components 5
 - factory default address 40
- storage node disk drive failure
 - troubleshooting 52
- storage server
 - cluster components 6
- supported
 - VLAN options 18
- supported components
 - color-coded Ethernet cables 13
 - grey Ethernet cables 13
- supported Fibre Channel adapters and switches
 - Brocade Fibre Channel switch 6
 - Fibre Channel expansion card 6
 - Fibre Channel SCSI/iSCSI dual channel adapter 6
 - Fibre Channel SCSI/iSCSI single channel adapter 6
 - Force 10 switch fabric module 6
 - SAN switch 2005 model H-16 6
- supported Gigabit switches
 - Cisco Catalyst Ethernet switch module 8
 - Cisco Catalyst Ethernet switch module 3750G-24T 8
 - Extreme Networks Summit 400 Gigabit Ethernet switch 8
 - Force 10 E600 Gigabit Ethernet switch 8
 - IBM eserver BladeCenter Gigabit Ethernet switch module 8
 - SMC Ethernet switch module 8624T 8
 - SMC Ethernet switch module 8648T 8
- switch cabling
 - 10/100/1000 Ethernet switch 25
 - Myrinet high-speed switch 25
- switching on
 - primary cabinets 30
- system board failures
 - troubleshooting 52
- system overview 1

T

- terminal server
 - cluster components 8
 - description 8
 - factory default address 40
- Topspin 12-port 4X InfiniBand line card
 - description 7
- Topspin 120 InfiniBand switch
 - description 7
- Topspin 120 switch
 - factory default address 40

- Topspin 270 InfiniBand management module
 - description 7
- Topspin 270 InfiniBand switch
 - description 7
- Topspin 270 InfiniBand switch fabric module
 - description 7
- Topspin InfiniBand host channel adapter
 - description 7
- Topspin InfiniBand host channel adapter PCI-Express
 - description 7
- Topspin InfiniBand host channel adapter PCI-X
 - description 7
- Topspin Infiniband host channel adapters
 - cluster components 7
- Topspin switch
 - cabling 26
- trademarks 68
- tree view
 - description 31
- troubleshooting
 - cluster node disk drive failure 52
 - storage node disk drive failure 52
 - system board failures 52
- turn on cluster
 - procedure 29
- turning on
 - expansion cabinets 30

U

- United States electronic emission Class A notice 71
- United States FCC Class A notice 71

V

- verifying
 - LCIT installation 31
- viewing
 - APC event log 59
 - CSM event log 59
 - LCIT event log 59
 - Linux event log 59
 - node event log 59
 - RSA event log 59
- VLAN options
 - supported 18
- Voltaire 9024 InfiniBand switch
 - description 7
- Voltaire switch
 - cabling 26

X

- xSeries 336 cluster node
 - factory default address 40
- xSeries 336 nodes
 - problem determination 51
- xSeries 336 problems
 - problem determination 53
- xSeries 346
 - management node 5

xSeries 346 nodes 51

Y

yellow CAT5E Ethernet cables
part number 13



Part Number: 25K8435

Printed in USA

(1P) P/N: 25K8435

