eServer Cluster 1350

Cluster 1350 Installation and Service

eServer Cluster 1350

Cluster 1350 Installation and Service

Note!

Before using this information and the product it supports, please review the information "Safety and environmental notices" on page xi.

June 2003 edition

This book is the current version of IBM eServer Cluster 1350 installation and service documentation.

IBM welcomes your comments. You may address your comments to the following address:

International Business Machines Corporation
Department 55JA, Mail Station P384
2455 South Road
Poughkeepsie, NY 12601-5400
United States of America

FAX (United States and Canada): 1+845+432-9405
FAX (Other Countries):
(Your international Access Code)+1+845+432+9405

IBMLink (United States customers only): IBMUSM10(MHVRCFS)
Internet e-mail: mhvrcfs@us.ibm.com

If you would like a reply, be sure to include your name, address, telephone number, or FAX number.

Make sure to include the following in your comment or note:
- Title of this book
- Page number or topic related to your comment

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© Copyright International Business Machines Corporation 2003. All rights reserved.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

© Copyright IBM Corp. 2003

Figures

© Copyright IBM Corp. 2003

Tables

© Copyright IBM Corp. 2003

Safety and environmental notices

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 notices

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

Safety Information

Before installing this product, read the Safety Information.

قبل تركيب هذا المنتج، يجب قراءة الملاحظات الآمنية

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

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

安裝本產品之前，請先閱讀「安全資訊」。

Prije instalacije ovog produkta obavezno pročitajte Sigurnosne Upute.

Před instalací tohoto produktu si přečtěte příručku bezpečnostních instrukcí.

Læs sikkerhedsforskrifterne, før du installerer dette produkt.

Lees voordat u dit product installeert eerst de veiligheidsvoorschriften.

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

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

Vor der Installation dieses Produkts die Sicherheitshinweise lesen.

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

לפני שתתקינו מוצר זה, קראו את הוראות הבטיחות.

A termék telepítése előtt olvassa el a Biztonsági előírásokat!

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

製品の設置の前に、安全情報をお読みください。

© Copyright IBM Corp. 2003

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 pečítaje Bezpečnostné predpisy.

Pred namestitvijo tega proizvoda preberite Varnostne informacije.

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

Läs säkerhetsinformationen innan du installerar den här produkten.

Important:

All caution and danger statements in this documentation begin with a number. This number is used to cross reference an English caution or danger statement with translated versions of the caution or danger statement in the *IBM NetBAY Rack Safety Information* book.

For example, if a caution statement begins with a number 1, translations for that caution statement appear in the *IBM NetBAY Rack Safety Information* book under statement 1.

Be sure to read all caution and danger statements in this documentation before performing the instructions. Read any additional safety information that comes with your server or optional device before you install the device.

Statement 2:

DANGER

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

Statement 3:

DANGER

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

Statement 4:

DANGER

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

To avoid a shock hazard:

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

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

Statement 7:

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.

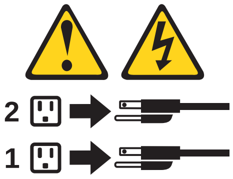

Statement 8:

DANGER

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

Statement 10:

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

- **Reduce the weight of the rack cabinet by removing equipment starting at the top of the rack cabinet. When possible, restore the rack cabinet to the configuration of the rack cabinet as you received it. If this configuration is not known, you must do the following:**
  - **Remove all devices in the 32U position and above.**
  - **Ensure that the heaviest devices are installed in the bottom of the rack cabinet.**
  - **Ensure that there are no empty U-levels between devices installed in the rack cabinet below the 32U level.**
- **If the rack cabinet you are relocating is part of a suite of rack cabinets, detach the rack cabinet from the suite.**
- **Inspect the route that you plan to take to eliminate potential hazards.**
- **Verify 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.**
- **Verify 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.

The following is a list of safety notices (in English only) pertaining to hardware maintenance tasks:

DANGER notices warn you of conditions or procedures that can result in death or severe personal injury.

CAUTION notices warn you of conditions or procedures that can cause personal injury that is neither lethal nor extremely hazardous.

ATTENTION emphasizes certain explanatory information and calls attention to statements in text, figures, and tables. Attention notices also highlight situations that could potentially cause damage to the equipment or loss of data.

Environmental notices

Product recycling and disposal

This product contains materials such as circuit boards, cables, electromagnetic compatibility gaskets, and connectors which might 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. You can find country-specific instructions at www.ibm.com/ibm/environment/products/prp.phtml.

This product might contain nickel-cadmium or lithium batteries in communication adapters. The batteries must be recycled or disposed of properly. Recycling facilities might not be available in your area. In the United States, IBM has established a collection process for reuse, recycling, or proper disposal of used sealed lead-acid, nickel-cadmium and nickel metal hydride batteries and battery packs from IBM equipment. For information on proper disposal of batteries in this product, please contact IBM at 1-800-426-4333. For information on disposal of batteries outside the United States, contact your local waste disposal or recycling facility.

Part 1. Introduction to Cluster 1350

© Copyright IBM Corp. 2003

Chapter 1. System overview

Contents

The Cluster 1350 can have a maximum of 512 nodes in addition to the one required Management Node. All nodes run the Linux operating system.

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. Expansion cabinets may contain Storage Nodes or mass-storage devices as well as computing nodes called Cluster Nodes. Expansion cabinets do not contain a Management Node or console. Figure 1 on page 4 shows an example of a Primary cabinet. Figure 2 on page 5 shows an example of an Expansion cabinet containing Cluster Nodes.

Figure 3 on page 6 shows an example of an Expansion cabinet containing storage controllers and mass storage.

© Copyright IBM Corp. 2003

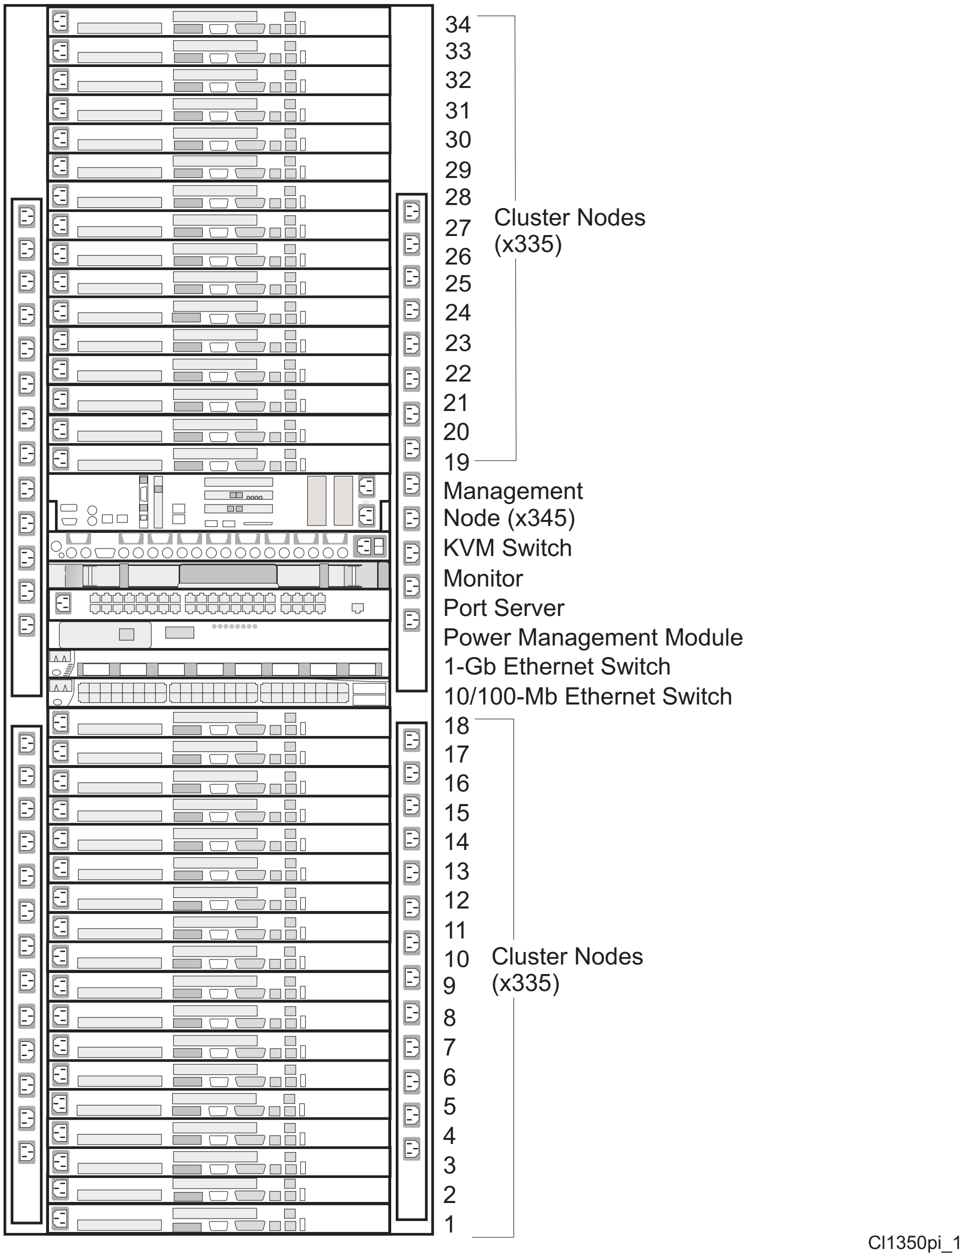

Figure 1. Example of an @server *Cluster 1350 Primary cabinet*

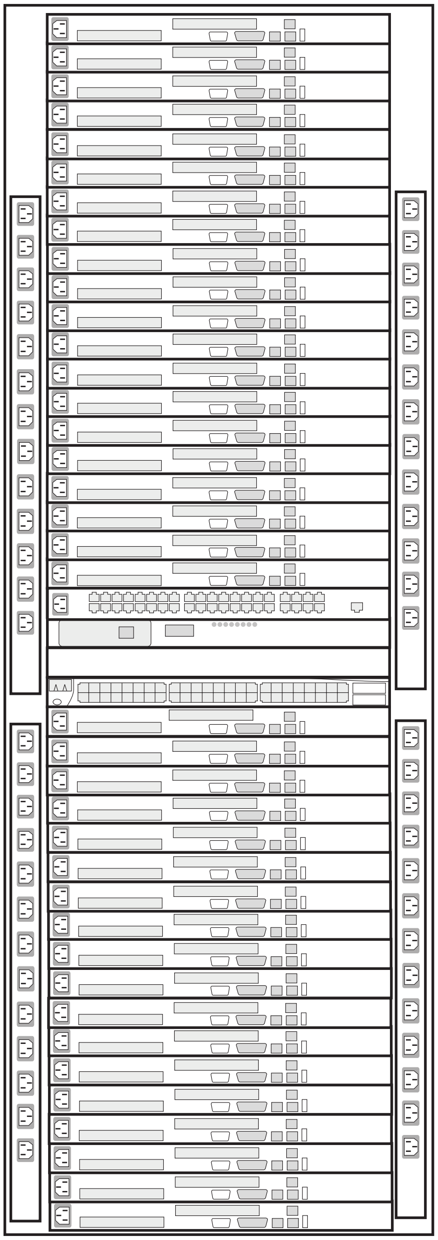

Cabinet 1	Cabinet 2	Cabinet 3	....	Cabinet 14	
38	76	114		Blank	Cluster Nodes (x335)
37	75	113		Blank	
36	74	112		Blank	
35	73	111		Blank	
34	72	110		Blank	
33	71	109		Blank	
32	70	108		Blank	
31	69	107		Blank	
30	68	106		Blank	
29	67	105		Blank	
28	66	104		Blank	
27	65	103		Blank	
26	64	102		Blank	
25	63	101		Blank	
24	62	100		Blank	
23	61	99		Blank	
22	60	98		Blank	
21	59	97		Blank	
20	58	96		Blank	
19	57	95		Blank	
Port Server					
Power Management Module					
1U Switch Option					
10/100-Mb Ethernet Switch					
18	56	94		512	Cluster Nodes (x335)
17	55	93		511	
16	54	92		510	
15	53	91		509	
14	52	90		508	
13	51	89		507	
12	50	88		506	
11	49	87		505	
10	48	86		504	
9	47	85		503	
8	46	84		502	
7	45	83		501	
6	44	82		500	
5	43	81		499	
4	42	80		498	
3	41	79		497	
2	40	78		496	
1	39	77		495	

CI1350pi-2

Figure 2. Example of an @server Cluster 1350 Expansion Cabinet with Cluster Nodes. This figure also shows how the node numbering scheme maps to other Expansion Cabinets.

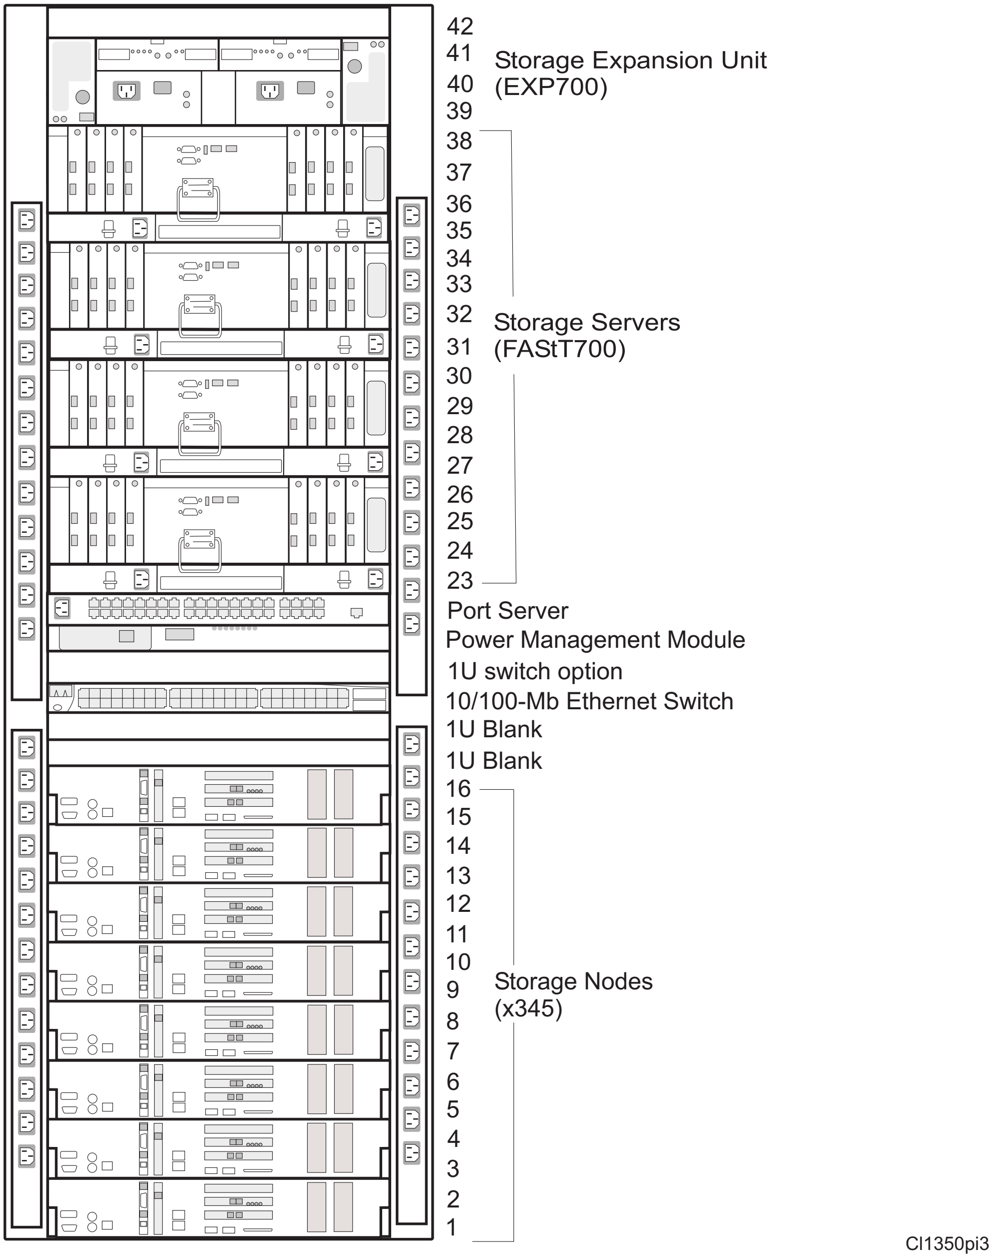

Figure 3. Example of an @server Cluster 1350 Expansion cabinet containing storage controllers and mass storage

The IBM® @server Cluster 1350 uses the following modules:

Cluster Node
The Cluster Nodes carry out the computational tasks in the cluster. The Cluster 1350 Cluster Node is either an IBM eServer xSeries™ 335 (x335) or IBM eServer BladeCenter running a supported Linux distribution. A cluster must contain at least four Cluster Nodes.

Management Node
Each cluster contains one Management Node, which provides system management for all modules in the cluster. The Cluster 1350 Management Node is an IBM eServer xSeries 345 (x345) running a supported Linux distribution.

Storage Node

The optional Storage Nodes manage the mass storage.

For tasks that do not require large amounts of mass storage, the Storage Node's onboard disk storage may be sufficient. The Cluster 1350 Storage Node is either an IBM eServer xSeries 345 (x345) or IBM eServer xSeries 360 (x360)..

More frequently, however, mass storage is desired for today's computing tasks. Each Storage Node can communicate with a FAStT200 Storage Server or FAStT700 Storage Server over a Fibre Channel connection.

Another mass storage solution is the IBM EXP300 SCSI expansion box. Each SCSI expansion box can be attached to any Cluster Node using a ServeRAID™ PCI card. The IBM EXP300 contains fourteen disk drives.

Storage Server

For this option the Cluster 1350 system uses the RAID-capable FAStT200 Storage Server or FAStT700 Storage Server. Each FAStT200 Storage Server adds up to ten internal drives to the storage capacity of the cluster. The FAStT700 Storage Server can support up to 224 drives contained in external expansion cabinets. Both enable the Storage Node to communicate with large RAID-protected arrays of storage.

Storage Expansion Unit

Each FAStT200 Storage Server in the cluster controls up to two IBM EXP500 Disk Storage Expansion Units, each of which expands the capacity of the Storage Server by ten disk drives. Each FAStT700 Storage Server in the cluster controls up to sixteen IBM EXP700 Disk Storage Expansion Units, each of which expands the capacity of the Storage Server by fourteen disk drives.

Another mass storage solution is the IBM EXP300 SCSI expansion box. Each SCSI expansion box can be attached to any Cluster Node using a ServeRAID PCI card. The IBM EXP300 contains fourteen disk drives.

Console

The console provides the monitor, keyboard, and mouse for the Management Node. The monitor is a fold-down flat-panel display that retracts into the rack.

KVM Switch

The KVM switch lets the console connect to all the different nodes in the cluster. The Cluster 1350 can use the IBM NetBAY 2x8 Console Switch or the IBM NetBAY Remote Console Manager to act as its KVM switch.

When using the IBM NetBAY 2x8 Console Switch, Storage and Management Nodes (x345s) connect directly to the switch. For Cluster Nodes (x335s) in the same rack, you can daisy-chain up to 40 nodes on one switch port.

When using the IBM NetBAY Remote Console Manager (RCM) up to 16 Storage and Management Nodes (x345s) can be daisy-chained to one port using CCO converters and CAT5 cables. For Cluster Nodes (x335s) in the same rack, you can daisy-chain up to 18 nodes on one switch port.

10/100-MB Ethernet Switch

The 10/100-MB Ethernet switch provides 10/100-MB Ethernet connections for the cluster. The Cluster 1350 uses the Cisco Ethernet switch Models 3550 XL (24-port) and 3550 XL (48-port). You can partition the switch to set up multiple independent LANs within the same switch.

Each model also provides two 1-GB Ethernet ports for communication with the Management Node or connection to the backbone of the network.

1-GB Ethernet Switch

The 1-GB Ethernet switch provides a 1-GB Ethernet trunk line between the Management Node and the Cluster and Storage Nodes. The Cluster 1350 uses the Cisco Ethernet switch Models 3508G (8-port). The 1-GB Ethernet switch uses either optical or copper cable with one of two optional GBICs.

Port Server

The port server provides serial connections for cluster modules. The Cluster 1350 uses the MRV® In-Reach® IR-8020–101 (20–port), or the MRV In-Reach IR-8040–101 (40–port).

The main purpose of the port server is to assign Ethernet addresses to cluster components. The port server can also act as a backup for Ethernet connections to download firmware and to check information stored in logs in cluster components.

High-Speed Myrinet Switch (optical)

This is an optional 2-GB switch for interconnecting Cluster Nodes and Storage Nodes. The Cluster 1350 uses the Myrinet® Models M3-E32 (5-slot), M3-E64 (9-slot), and M3-E128 (17). The high-speed switch can replace the optional secondary Ethernet. It requires a Myrinet PCI adapter in each Cluster Node and Storage Node. The Myrinet Switch uses an optical cable.

High-Speed Cisco Ethernet Switch (copper)

The Cluster 1350 can also use the Cisco® Catalyst® 4003 (3–slot) and Cisco Catalyst 4006 (6–slot) switches with optional 10/100/1000 line cards to provide a high-speed Cluster VLAN or to provide a second high-speed Ethernet network at a lower cost than the high-performance Myrinet network.

Power Management Module

The Power Management Module provides power to the service processors (RSA boards) and to the port servers. The Cluster 1350 uses the APC® MasterSwitch™ Model AP9212. The Power Management Module can supply up to eight connections. It provides the ability to power-cycle a component remotely.

Power Distribution Unit (PDU)

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

- Rack PDU
- Front end PDU

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

DANGER

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

Related Topics

You can also refer to the following information:

- **IBM:**

  x335

  xSeries 335 Installation Guide:
  ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/33p2612.pdf
  xSeries 335 User's Guide:
  ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/33p2611.pdf
  xSeries 335 Hardware Maintenance Manual and Troubleshooting Guide:
  ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/48p9908.pdf

  BladeCenter

  eServer BladeCenter Planning and Installation Guide:
  ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/ga27-4327-00.pdf
  eServer BladeCenter Hardware Maintenance Manual and Troubleshooting Guide:
  ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/71p9885.pdf

  x345

  xSeries 345 Installation Guide:
  ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/48p9726.pdf
  xSeries 345 User's Guide:
  ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/48p9717.pdf
  xSeries 345 Hardware Maintenance Manual and Troubleshooting Guide:
  ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/48p9718.pdf

  x360

  xSeries 360 Installation Guide:
  ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/48p9794.pdf
  xSeries 360 User's Reference:
  ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/48p9793.pdf
  xSeries 360 Hardware Maintenance Manual:
  ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/24p2967.pdf

  Storage Server

  FAStT Host Adapter Installation and User's Guide:
  ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers/25p1663.pdf
  Fibre Array Storage Technology, a FAStT Introduction(Redbook):
  http://www.redbooks.ibm.com/pubs/pdfs/redbooks/sg246246.pdf
  FAStT 700 Installation Guide:
  ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/32p0171.pdf
  FAStT 700EXP Installation and User's Guide:
  ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/32p0178.pdf

  Monitor

IBM NetBAY 1U Flat Panel Monitor Console Kit Installation and Maintenance Guide:
ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/02r2712.pdf

- **MRV:**

  In-Reach 8000 Series Port server

  Product information:
  http://service.mrv.com/support/index.cfm
- **APC:**

  Power Management Module

  Product information:
  http://www.apcc.com/resource/include/techspec_index.cfm?base_sku=AP9212
  MasterSwitch Power Distribution Unit User's Guide:
  http://sturgeon.apcc.com/techref.nsf/partnum/990-6018e
  Troubleshooting information:
  http://www.apcc.com/support/kbase.cfm
- **Cisco:**

  10/100-MB and 1-GB Ethernet switches

  General Product information:
  http://www.cisco.com/univercd/cc/td/doc/product/lan/
  Catalyst 3550 Series XL Hardware Installation Guide:
  http://www.cisco.com/univercd/cc/td/doc/product/lan/c3550/1219ea1/3550hig/index.htm
  Quick Start Guide Catalyst 3550 Series XL Switches:
  http://www.cisco.com/univercd/cc/td/doc/product/lan/c2900xl/29_35xu/3500.htm
  Quick Start Product information for Catalyst 4000 Series switches:
  http://www.cisco.com/univercd/cc/td/doc/product/lan/cat4000/index.htm
- **Myrinet:**

  High-speed switches

  Myrinet software and documentation:
  http://www.myri.com/scs/index.html

Part 2. Installation

© Copyright IBM Corp. 2003

Chapter 2. Unpacking the eServer Cluster 1350

Contents

The eServer Cluster 1350 is designed to be set up by IBM. The customer should complete the following checklist before IBM arrives on site to finish the installation:

__ 1. Review the legal and safety information.

__ 2. Review the physical, environmental, and electrical requirements for the Cluster 1350 as outlined in **Cluster 1350 Preinstallation Planning** and ensure the final installation site is ready.

__ 3. Unpack the cabinet(s) only, but not the other boxes. Depending on the configuration of the cluster you ordered, some equipment may have been removed from the racks to satisfy shipping requirements.

ATTENTION!

Do not attempt to replace any equipment that was removed from the racks. IBM will install all equipment back into its proper location as part of the normal setup process.

__ 4. Using the order placed for the system, identify the Primary cabinet and verify its contents. If equipment was removed prior to shipping, check the Bill of Materials to ensure that all the equipment required for the Primary cabinet was shipped with the order.

__ 5. Using the order placed for the system, identify the Expansion cabinets and verify their contents. If equipment was removed prior to shipping, check the Bill of Materials to ensure that all the equipment required for the Expansion cabinets was shipped with the order.

__ 6. Dispose of the packing material that came with the cabinets.

__ 7. Move the cabinets and any boxes containing extra equipment or other material to the installation site. IBM will take care of final cabinet placement.

__ 8. Arrange for a phone line near the cabinets.

The IBM Service team will perform the final cabling and installation steps. After the system is completely installed, you can connect your network cables.

© Copyright IBM Corp. 2003

Chapter 3. Placing the cabinets

Contents

Customer responsibilities

Once the contents of all the cabinets of the Cluster 1350 are verified, the customer should move the cabinets and any boxes containing extra equipment and other materials to the location they have prepared for the installation.

Physical, environmental, and electrical requirements were outlined in the **IBM eServer Cluster 1350 Preinstallation Planning** manual. Do not move the cabinets to their final location if proper preparations are still being made to the space.

The Cluster 1350 was manufactured according to information provided at the time the order was placed. The side-to-side and front-to-back clearances for each cabinet are directly related to the load carrying capability of the floor in the location of the installation. Cabinet to cabinet cabling harnesses were custom made for each order to allow for proper spacing of the cabinets. If the location of the installation has changed since the time the order was placed the customer should review the physical, environmental, and electrical requirements outlined in the preinstallation manual to ensure there are no incompatibilities.

Using the packing slip enclosed with the Cluster 1350, the customer should place the cabinets in their approximate final location. Each cabinet is labeled to help in this process.

The IBM Service team will take care of final cabinet placement and perform the rest of the cabling and installation steps.

Installer responsibilities

IBM Global Services or an IBM Authorized Business Partner will perform the remaining activities in the installation process.

Cabinet placement

Use the following guidelines when placing cabinets:

- Cabinets can be placed side-by-side in contact with one another. Remember that to service any Power Distribution Unit (PDU) in the cabinet you will have to remove the side covers. At least 30 inches of *working clearance* is required to allow removal of a side cover and permit access to the PDU. If the cabinets are placed side-by-side in contact with one another you should have enough extra space around the cluster to allow for movement of the cabinets in the event 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 were custom made to fit the order. If the location of the installation has changed since the time the order was placed, review the physical, environmental, and electrical requirements outlined in the preinstallation manual to ensure there are no incompatibilities.

© Copyright IBM Corp. 2003

- Check that the cabinets are arranged properly and adjust if needed. Refer to the packing slip and the cabinet labels to verify that all cabinets are in their proper location.

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.

Take the following steps to finish placing the cabinets:

__ 1. Inspect the cabinets, components, and cable connections for shipping damage.

__ 2. Install the cabinet stabilizer on each cabinet as required.
   - Refer to Figure 4 for stabilizer kit installation

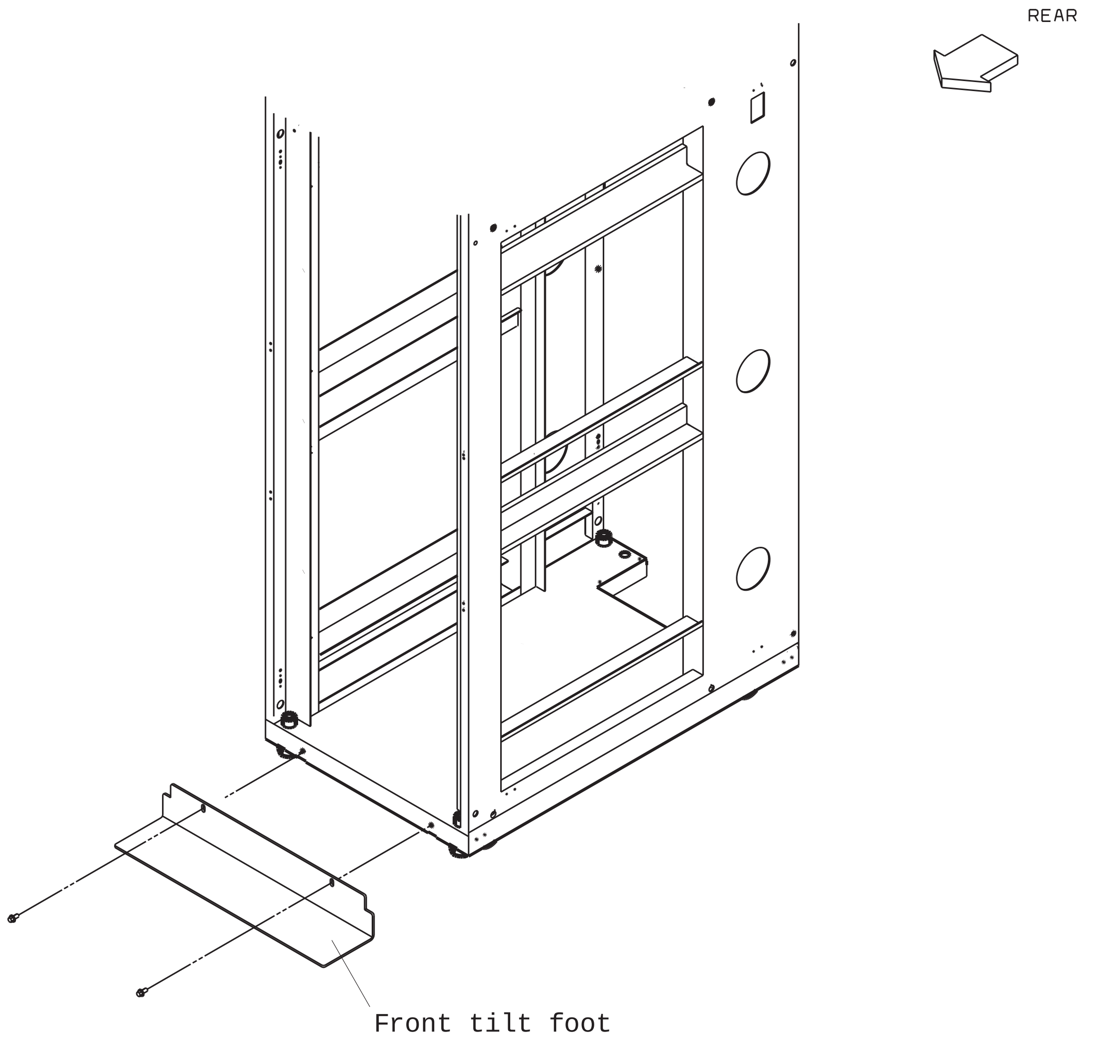

Figure 4. Frame stabilizer (tilt) foot installation

Chapter 4. Cabling

Contents

Overview of Cluster 1350 cabling

The various types of cables in the Cluster 1350 system perform the following functions:

Management VLAN
Provides the private virtual LAN (VLAN) to manage the components in the cluster. This VLAN includes the following connections:

- RS485 connections to all Cluster Nodes and Storage Nodes through the RSA boards. These enable diagnostics and monitoring for the Cluster and Storage Nodes.
- Serial connections to all cluster components. These provide a path for configuration of components in the cluster.
- 10/100 Ethernet connections from the RSA card in the Management Node to the 10/100 Ethernet Switch.

Primary Cluster VLAN
Provides a 10/100 or a 10/100/1000 Ethernet (depending on the VLAN type selected) for communications with Cluster Nodes and Storage Nodes. This VLAN includes the following connections:

- A 10/100 or 10/100/1000 Ethernet 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 communications within the cluster.

Optional Secondary Cluster VLAN
Provides a second 10/100/1000 Ethernet or 2–GB Myrinet switch for communications with Cluster and Storage Nodes. There are several options for the secondary cluster VLAN:

- A 10/100/1000 Ethernet using a Cisco 4003 or 4006 switch
- A 2-GB Ethernet using a high-speed Myrinet switch. The Myrinet switch uses optical cables for communication with Cluster and Storage Nodes.

KVM Connects the KVM ports on all nodes (Cluster, Storage, and Management) to a single console through a central switch.

Fibre Channel
Provides Fibre Channel connections between the Storage Nodes and the Storage Servers, and between the Storage Servers and the Storage Expansion Units.

Power Provides the power to the cluster components. This includes both the AC power provided to the entire cabinet through the PDUs and remote power provided to the RSA boards and the port servers through the Power Management Module.

VLAN options

The Cluster 1350 supports a variety of VLAN options. Currently, there are six basic configurations. Point-to-point wiring information is printed on each cable. Check

© Copyright IBM Corp. 2003

the information on the cables in the Primary rack and refer to the following tables to determine which VLAN option was used in the cluster.

Table 1. Type 1 VLAN. 10/100 Ethernet

Device	**Management VLAN**	**10/100 Primary Cluster VLAN**	**Comments**
Management Node	Ethernet 2 connects to Cisco 3550	Ethernet 1 connects to Cisco 3508 GBIC	
KVM switch	Connects to Cisco 3550		
Cisco 3508 Gbit Switch		GBIC connects to Management Node Ethernet 1	
iTouch port server	Conncets toCisco 3550		
APC switch	Connects to Cisco 3550		
Cisco 3550 10/100 switch		GBIC connects to Cisco 3508	
Cluster Nodes		Ethernet 0 connects to Cisco 3550	
Storage Nodes	Ethernet 2 connects to Cisco 3550	Ethernet 1 connects to Cisco 3508 GBIC	
FAStT200HA	Connects to Cisco 3550		Uses both jacks
FAStT700	Connects to Cisco 3550		Uses both jacks

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

Device	**Management VLAN**	**Gbit Primary Cluster VLAN**	**Comments**
Management Node	• Ethernet 1 connects to Cisco 400x • Ethernet 2 connects to Sup I or Sup III		06P3701 or 22P7801
KVM switch	Connects to Cisco 400x		
iTouch port server	Connects to Cisco 400x		
APC switch	Connects to Cisco 400x		
Cisco 4003 and/or 4006 switch	• Sup I connects to Management Node • Sup III connects to Management Node	• Gbit connects to Management Node Ethernet1 • Sup III uplink connects to Management Node PCI card	
Cluster Nodes		Ethernet 0 connects to Cisco 400x	

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

Storage Nodes	Ethernet 2 connects to Cisco 400x	Ethernet 1 connects to Cisco 400x	
FAStT200HA	Connects to Cisco 400x		Uses both jacks
FAStT700	Connects to Cisco 400x		Uses both jacks

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

Device	**Management VLAN**	**10/100 Primary Cluster VLAN**	**Gbit customer public high speed VLAN**
Management Node	Ethernet 2 connects to Cisco 3550	Ethernet 1 connects to Cisco 3508 GBIC	
KVM switch	Connects to Cisco 3550		
Cisco 3508 (4003) Cisco 3508 (4006)		Copper GBIC connects to Management Node Ethernet 1 Copper GBIC connects to Management Node Ethernet 1	Copper GBIC connects to Cisco 4003 Gbit Fiber GBIC connects to Sup III uplink
iTouch port server	Connects to Cisco 3550		
APC switch	Connects to Cisco 3550		
Cisco 3550 10/100 switch		Copper GBIC connects to Cisco 3508	
Cisco 4003 and/or 4006 switch	SupI connects to Cisco 3550 SupIII connects to Cisco 3550		Gbit connects to 3508 copper GBIC SupIII uplink connects to Cisco 3508 fiber GBIC
Cluster Nodes		Ethernet 0 connects to Cisco 3550	Ethernet 2 connects to Cisco 400x Gbit
Storage Nodes	Ethernet 1 Alias connects to Cisco 3508 copper GBIC	Ethernet 1 Alias connects to Cisco 3508 copper GBIC	Ethernet 2 connects to Cisco 400x Gbit
FAStT200HA	Connects to Cisco 3550		
FAStT700	Connects to Cisco 3550		

Table 4. 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**

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

Management Node	Ethernet 2 connects to Cisco 3550	Ethernet 1 connects to Cisco 3508 copper GBIC	
KVM switch	Connects to Cisco 3550		
3508 Gbit switch		Copper GBIC connects to Management Node Ethernet 1	
iTouch port server	Connects to Cisco 3550		
APC switch	Connects to Cisco 3550		
Cisco 3550 10/100 switch		Copper GBIC connects to Cisco 3508	
Myrinet 32, 64, or 128 (both jacks)	Connects to Cisco 3550		
Cluster Nodes		Ethernet 0 connects to Cisco 3550	Myrinet adapter
Storage Nodes	Ethernet 2 connects to Cisco 3550	Ethernet 1 connects to copper Cisco 3508 GBIC	Myrinet adapter
FAStT200HA	Connects to Cisco 3550		
FAStT700	Connects to Cisco 3550		

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

Device	**Management VLAN**	**Gbit Primary Cluster VLAN**	**Gbit customer public high speed VLAN**
Management Node	• Ethernet 1 Alias connects to Cisco 400x • Ethernet 2 connects to Sup I or Sup III	• Ethernet 1 connects to 4003 • Fiber Channel PCI card connects to 4006 SupIII uplink	Copper or Fiber Channel PCI connects to public network
KVM switch	Connects to Cisco 400x		
iTouch port server	Connects to Cisco 400x		
APC switch	Connects to Cisco 400x		

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

Cisco 4003 and/or 4006 switch	• Sup I connects to Management Node Ethernet 2 • Sup III connects to Management Node Ethernet 2	• Gbit connects to Management Node Ethernet 1 • Sup III uplink 1 connects to Fiber Channel PCI card	SupIII 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
FAStT200HA (both jacks)	Connects to Cisco 400x		
FAStT700 (both jacks)	Connects to Cisco 400x		

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

Device	**Management VLAN**	**Gbit Primary Cluster VLAN**	**Myrinet customer public high speed VLAN**
Management Node	• Ethernet 1 Alias connects to Cisco 400x • Ethernet 2 connects to Sup I or Sup III	• Ethernet 1 Alias connects to Cisco 400x • Fiber Channel PCI card (06P3701 or 22P78021) connects to Cisco 4006 Sup III uplink	
KVM switch	Connects to Cisco 400x		
iTouch port server	Connects to Cisco 400x		
APC switch	Connects to Cisco 400x		
Myrinet 32, 64, or 128 (both jacks)	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
FAStT200HA	Connects to Cisco 400x		
FAStT700	Connects to Cisco 400x		

For large clusters that use VLAN types 2, 5 or 6 it is necessary to add additional Cisco 400x switches. If the cluster you are working on has more than one Cisco 400x switch in the Primary rack then refer to the following tables.

Table 7. Type 2 VLAN with multiple Cisco 400x switches

Device	**Management VLAN**	**Gbit Primary Cluster VLAN**	**Comments**
Management Node	Ethernet 2 connects to Cisco 3550	Ethernet 1 connects to Cisco 3508 copper GBIC	
KVM switch	Connects to Cisco 3550		
iTouch port server	Connects to Cisco 3550		
3508 Gbit switch		Connects to 3550 copper GBIC	
APC switch	Connects to Cisco 3550		
3550 10/100 switch		Cisco 3550 copper uplink to Cisco 3508	
Cisco 4003 and/or 4006 switch	• Sup I connects to 3550 • Sup III connects to 3550	• Gbit connects to 3508 copper GBIC • Sup III uplink connects to 3508 fiber GBIC	
Cluster Nodes		Ethernet 0 connects to Cisco 400x	
Storage Nodes	Ethernet 2 connects to Cisco 3550	Ethernet 1 connects to Cisco 400x	
FAStT200HA	Connects to Cisco 3550		Uses both jacks
FAStT700	Connects to Cisco 3550		Uses both jacks

Table 8. Type 5 VLAN with multiple Cisco 400x switches

Device	**Management VLAN**	**Gbit Primary Cluster VLAN**	**Gbit customer public high speed VLAN**
Management Node	Ethernet 2 connects to Cisco 3550	Ethernet 1 connects to Cisco 3508 copper GBIC	Fiber Channel PCI connects to public network
KVM switch	Connects to Cisco 3550		
iTouch port server	Connects to Cisco 3550		
3508 Gbit switch	Connects to Cisco 3550 copper uplink	Connects to Management Node Ethernet 1	Fiber GBIC connects to Management Node
APC switch	Connects to Cisco 3550		
3550 10/100 switch	Connects to Management Node Ethernet 2		

Table 8. Type 5 VLAN with multiple Cisco 400x switches (continued)

Cisco 4003 and/or 4006 switch	• Sup I connects to 3550 • Sup III connects to 3550	• Gbit connects to 3508 copper GBIC • Sup III uplink 1 connects to 3508 fiber GBIC	• Gbit connects to 3508 copper GBIC • Sup III uplink 2 connects to 3508
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
FAStT200HA (both jacks)	Connects to Cisco 3550		
FAStT700 (both jacks)	Connects to Cisco 3550		

Table 9. 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 3550	Ethernet 1 connects to Cisco 3508 copper GBIC	Fiber Channel PCI connects to public network
KVM switch	Connects to Cisco 3550		
iTouch port server	Connects to Cisco 3550		
3508 Gbit switch	Connects to Cisco 3550 GBIC	Connects to Management Node Ethernet 1	
APC switch	Connects to Cisco 3550		
3550 10/100 switch	Connects to Management Node Ethernet 2		
Cisco 4003 and/or 4006 switch	• Sup I connects to Cisco 3550 • Sup III connects to Cisco 3550	• Gbit connects to Cisco 3508 copper GBIC • Sup III uplink connects to Cisco 3508 fiber GBIC	
Myrinet 32, 64, or 128 (both jacks)	Connects to Cisco 3550		
Cluster Nodes		Ethernet 0 connects to Cisco 400x	Myrinet adapter
Storage Nodes	Ethernet 2 connects to Cisco 3550	Ethernet 1 connects to Cisco 400x	Myrinet adapter
FAStT200HA	Connects to Cisco 3550		
FAStT700	Connects to Cisco 3550		

General information

Most of the cabling in a Cluster 1350 system will be installed during manufacturing. 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 will need to be reconnected at the customer site.

Attaching 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 will identify the device or node it connects to, and where appropriate, its port number.

Depending on the country of manufacture the label scheme will vary. Before you begin attaching cables take some time to familiarize yourself with the information on the labels.

When initially installing a Cluster 1350 start with the Primary cabinet. Once you have attached the intracabinet cables inside the Primary cabinet, move on to each Expansion cabinet and use the information printed on each cable label to properly attach the cables found there.

Once you have reattached any cables in the Primary cabinet and Expansion cabinets you can move on to attaching the cables that run between cabinets. This is called the intercabinet cabling and the following types of cables are involved:

- 1 or 2 GB Fibre Channel (optical)
- 1 GB Ethernet (optical)
- 2 GB Myrinet (optical)
- 10/100/1000 Ethernet (copper)
- KVM (copper)

For a complete listing of all available cables and their part numbers refer to the Cluster 1350 information contained on the IBM InfoTips web site.The same information is also available in the IBM Current Object Repository (CORE) system.

1-GB Ethernet cabling

The 1-GB Ethernet provides a high-speed optical trunk line for VLAN communications with the Management Node. Figure 5 on page 25 schematically shows a possible intercabinet cabling for a large cluster configuration. VLAN types 1, 3, and 4 would follow this model.

ATTENTION!

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

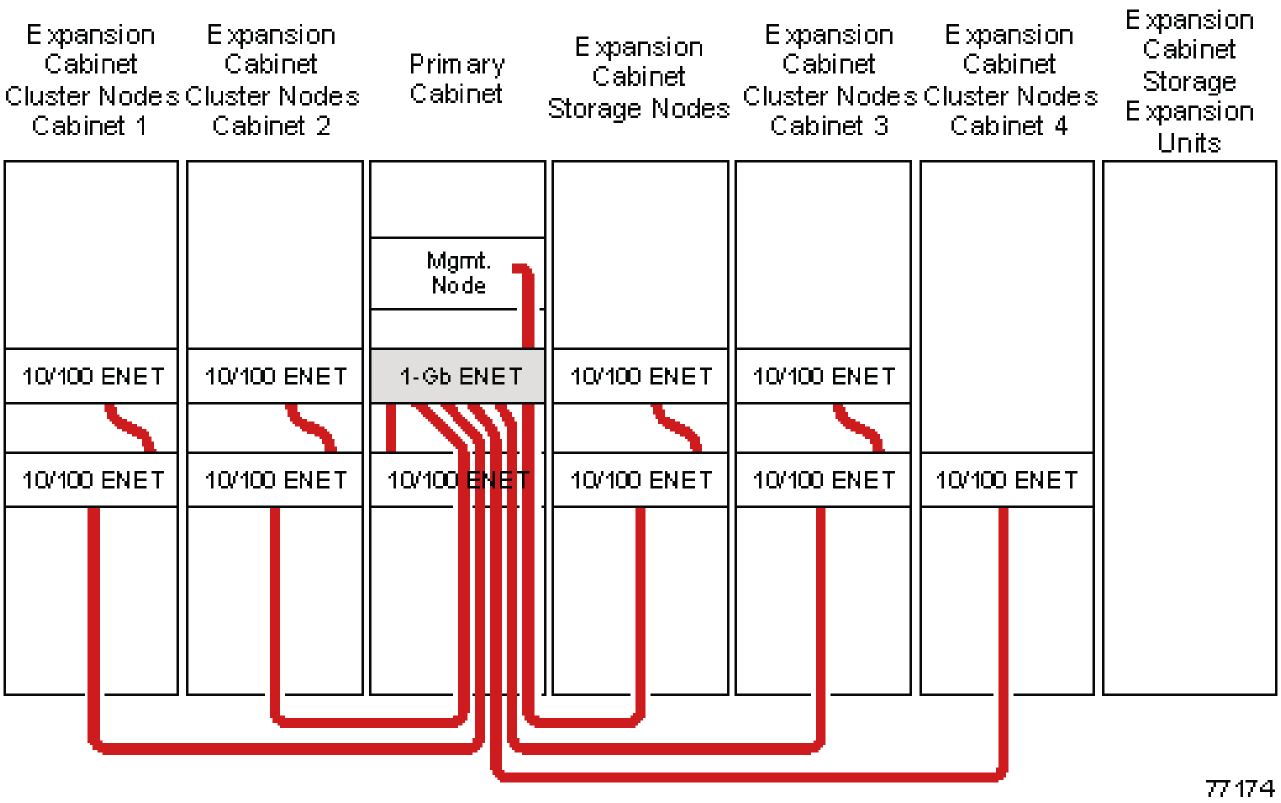
Expansion Cabinet Cluster Nodes Cabinet 1
Expansion Cabinet Cluster Nodes Cabinet 2
Primary Cabinet
Expansion Cabinet Storage Nodes
Expansion Cabinet Cluster Nodes Cabinet 3
Expansion Cabinet Cluster Nodes Cabinet 4
Expansion Cabinet Storage Expansion Units
Mgmt. Node
10/100 ENET
1-Gb ENET
77174

Figure 5. Intercabinet cabling for the 1-GB Ethernet connections (VLAN types 1, 3 and 4)

High-speed (Myrinet) switch cabling

The Myrinet high-speed switch provides an optional 2-GB optical network for communications between Cluster Nodes and Storage Nodes. Figure 6 on page 26 shows a schematic of the Myrinet optical cabling in a large cluster. VLAN types 2, 5, and 6 would follow this model.

ATTENTION!

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

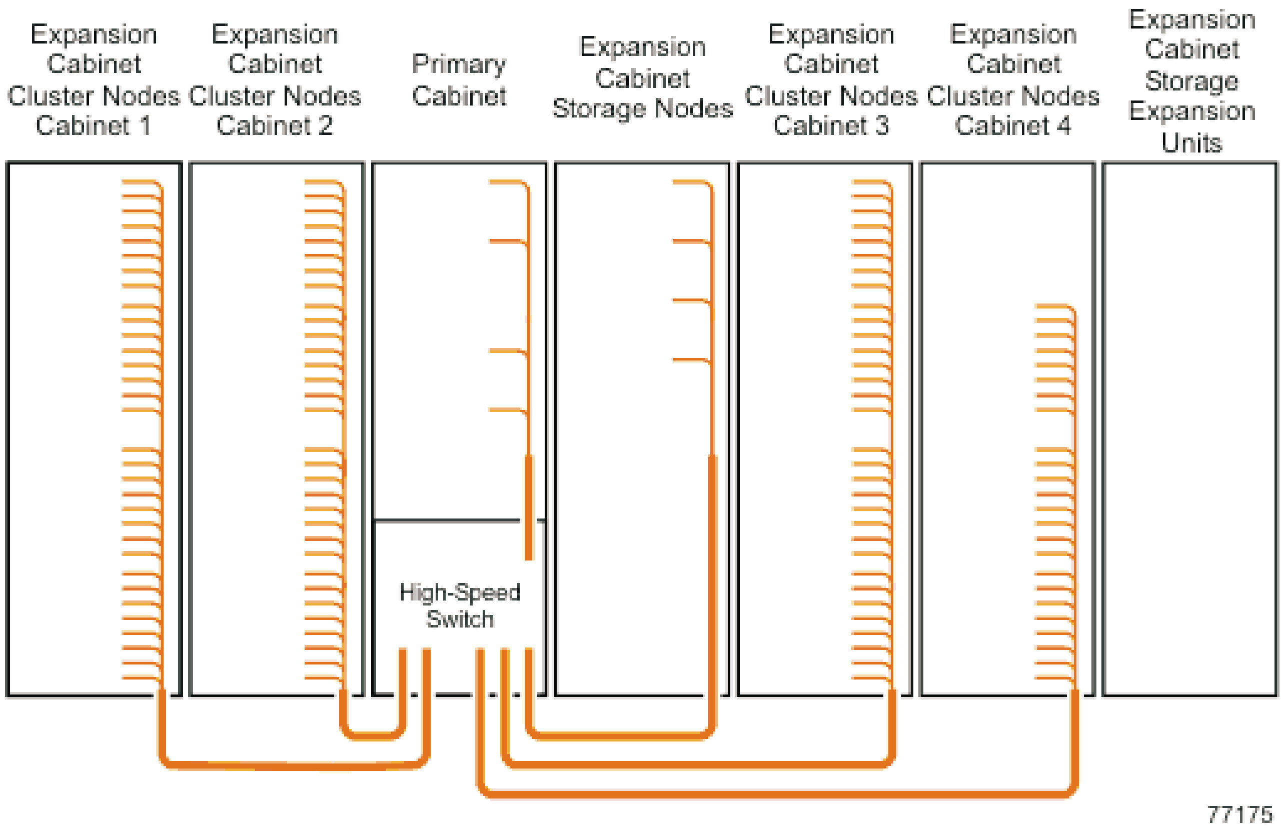
Expansion Cabinet Cluster Nodes Cabinet 1
Expansion Cabinet Cluster Nodes Cabinet 2
Primary Cabinet
Expansion Cabinet Storage Nodes
Expansion Cabinet Cluster Nodes Cabinet 3
Expansion Cabinet Cluster Nodes Cabinet 4
Expansion Cabinet Storage Expansion Units
High-Speed Switch
77175

Figure 6. Intercabinet cabling for Myrinet connections (VLAN types 4 and 6)

10/100/1000 Ethernet cabling

The 10/100/1000 Ehternet switch provides an optional 10/100/1000 network for communications between Cluster Nodes and Storage Nodes. Figure 6 shows a schematic of the cabling in a large cluster. VLAN types 2, 3,5 and 6 would follow this model.

ATTENTION!

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

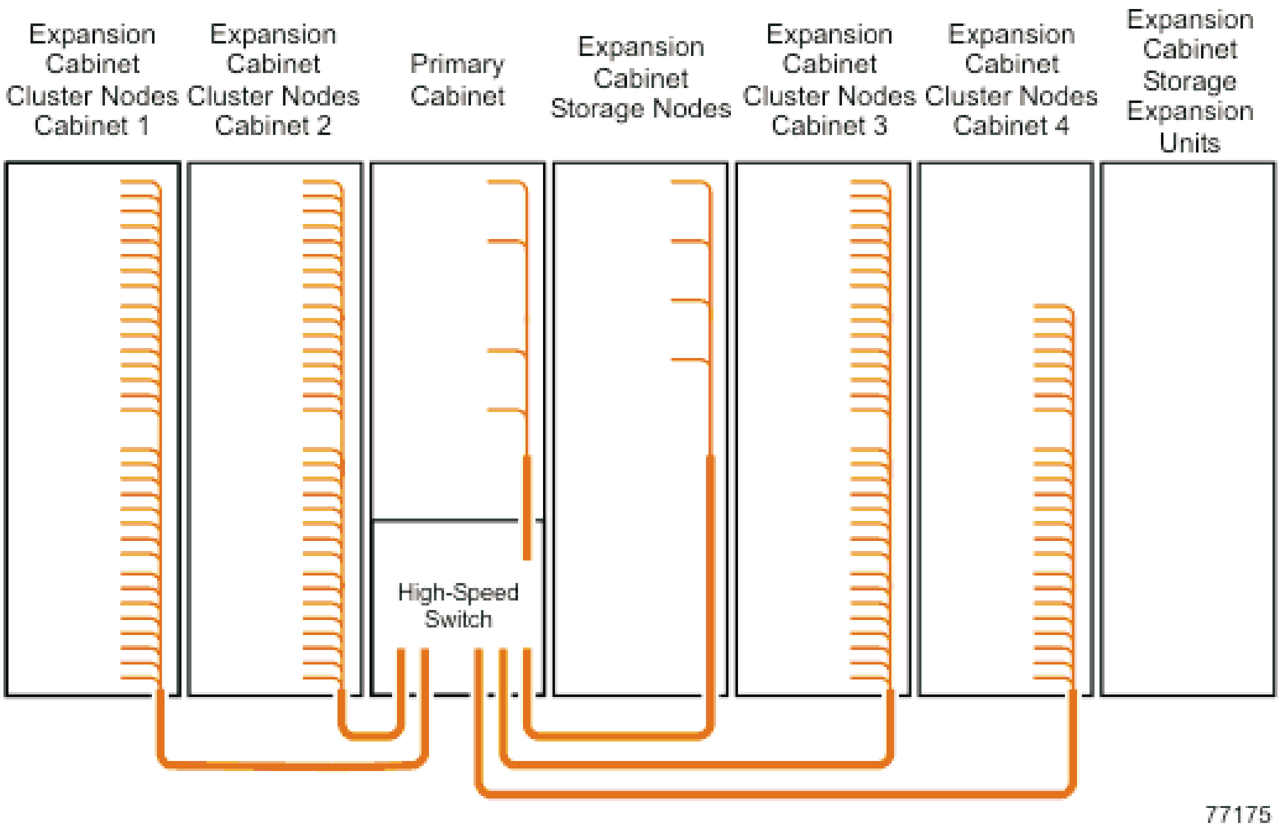
Expansion Cabinet Cluster Nodes Cabinet 1
Expansion Cabinet Cluster Nodes Cabinet 2
Primary Cabinet
Expansion Cabinet Storage Nodes
Expansion Cabinet Cluster Nodes Cabinet 3
Expansion Cabinet Cluster Nodes Cabinet 4
Expansion Cabinet Storage Expansion Units
High-Speed Switch
77175

Figure 7. Intercabinet cabling for Gbit Ethernet connections (VLAN types 2,3,5 and 6)

Fibre Channel cabling

Fibre Channel is used to connect Storage Nodes to Storage Servers , and to connect Storage Servers to Storage Expansion Units . Figure 8 on page 28 shows a schematic diagram of Fibre Channel Cabling in a large cluster.

ATTENTION!

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

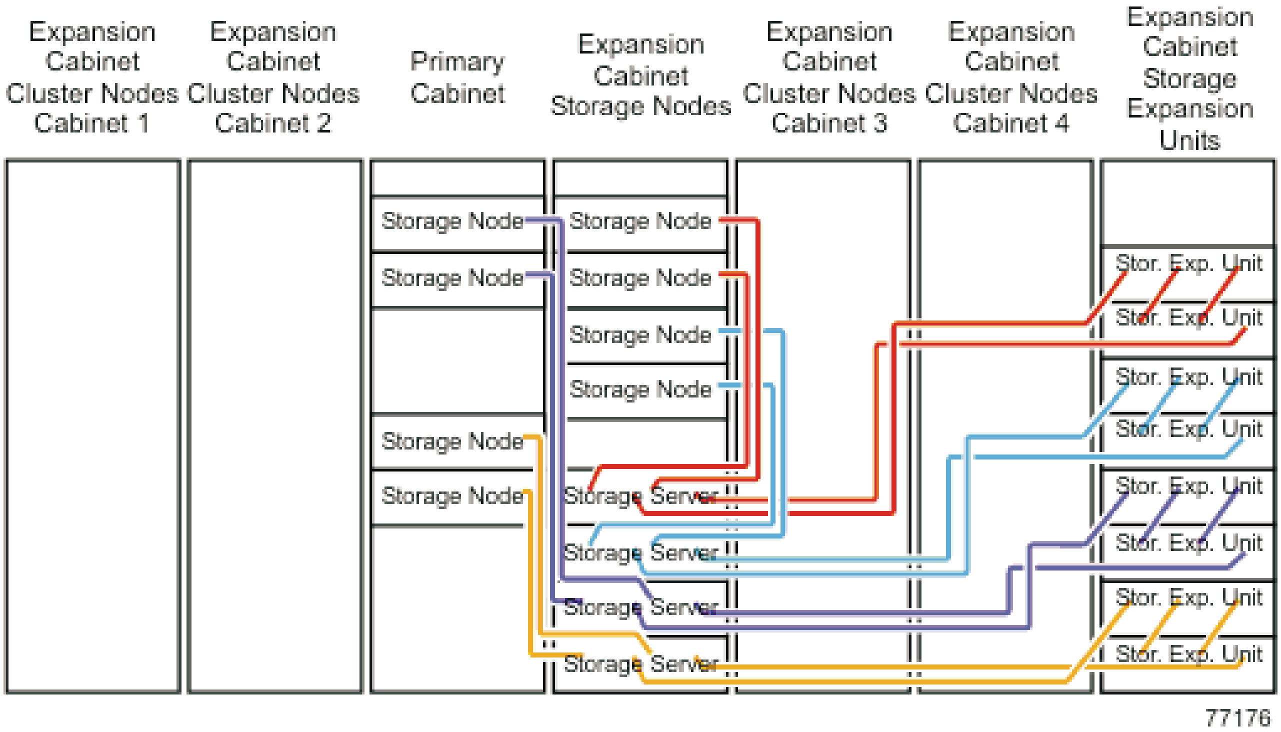

Figure 8. Intercabinet cabling for Fibre Channel connections.. To avoid making the Storage Node a single point of failure, connect two Storage Nodes to each Storage Server, as shown.

KVM cabling

The KVM switch allows a maximum of eight connections . Figure 9 on page 29 shows an example of KVM cabling for a cluster configuration. 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 daisy-chained.

  Cluster Nodes (x335s) can be daisy-chained onto 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 may require a second KVM switch. The second switch should be located 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.

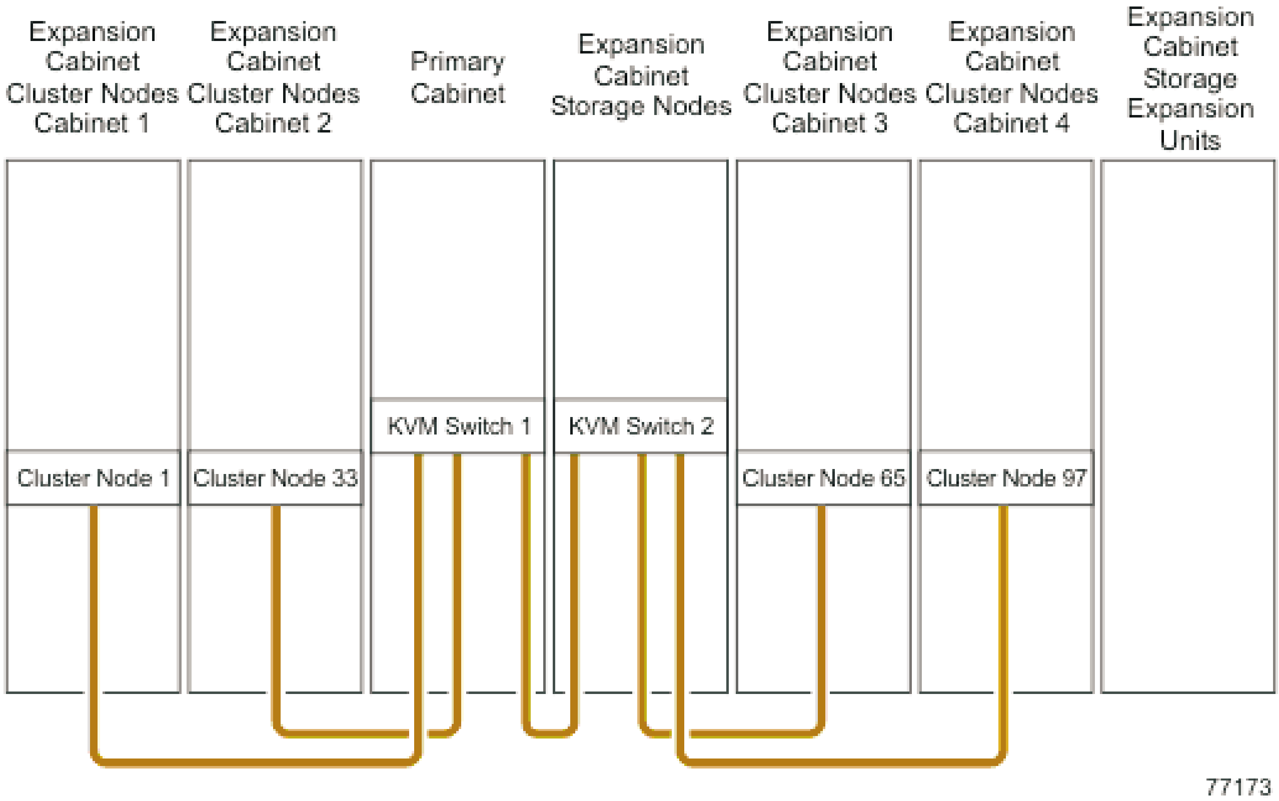
Expansion Cabinet Cluster Nodes Cabinet 1
Expansion Cabinet Cluster Nodes Cabinet 2
Primary Cabinet
Expansion Cabinet Storage Nodes
Expansion Cabinet Cluster Nodes Cabinet 3
Expansion Cabinet Cluster Nodes Cabinet 4
Expansion Cabinet Storage Expansion Units
KVM Switch 1
KVM Switch 2
Cluster Node 1
Cluster Node 33
Cluster Node 65
Cluster Node 97
77173

Figure 9. Intercabinet cabling for KVM connections

Remote Console Manager (RCM) cabling

The Remote Console Manager (RCM) switch has sixteen ACT connections (KVM over RJ45/CAT5) along with 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 daisy-chained.

  Up to 40 Cluster Nodes (x335s) can be daisy-chained to each ACT port on the RCM. The Management Node and all the Storage Nodes may also be daisy-chained, with up to sixteen per ACT port. Multiple RCMs may not be daisy-chained together. The RCM may be connected to an Ethernet network to allow for remote access to the consoles of the servers over the network..

Replacing a defective cable in a harness

If a cable in a harness is defective, replace the cable as follows:

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. If that can not be done, 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 accidentally been left unconnected.

3. Install a single cable between the two empty ports. Use a wire tie to attach the cable to the harness that contains the defective cable. This identifies the replacement cable as belonging to this harness.
4. Label the replacement cable so it is clearly identified as a replacement.

Chapter 5. Power up the cluster

Contents

Initial Cluster 1350 power-on procedure

The IBM eServer Cluster 1350 is shipped without an operating system installed. Before initially powering-on an entire Cluster 1350 system, first check all the connections in the Expansion cabinets and Primary cabinet. Once you have verified all connections are secure, power-on the Expansion cabinets containing Storage Nodes, Storage Servers and Storage Expansion Units. Power-on the Primary cabinet last.

Checking connections in the Expansion cabinets

1. Verify that the breaker switches for the customer's source power are all turned off.
2. Open the side and rear doors of the cabinet.
3. From the side of the cabinet check that all the power cables between the rack PDUs and the front-end PDUs are fully seated.
4. From the back of the cabinet, push on all the power plugs running from the rack-mounted devices to the PDUs to verify that the cables are fully seated.
5. Connect power to the PDUs.
   a. Plug the power cable into the PDU.
   b. Draw the power cord through the opening at the base of the cabinet.
   c. Plug the power cable into the wall outlet or other appropriate receptacle.
   d. Turn on the breaker switch for the customer's source power.
   e. Ensure the PDU circuit breakers are turned on.
6. Verify that all internal PDUs are powered up by viewing the power-on LEDs on components that are connected to the PDUs.
   - Servers display a blinking green light on the front panel when power is applied.
   - The following devices have no power switch and will power up automatically when the PDUs are powered up. Verify that these components have power applied.
     - KVM Switch
     - Cisco 10/100 Switch
     - Cisco Gigabit Switch
     - In-Reach Port Server

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

Checking connections in the Primary cabinet

1. Verify that the breaker switches for the customer's source power are all turned off.
2. Open the side and rear doors of the cabinet.
3. From the side of the cabinet check that all the power cables between the rack PDUs and the front-end PDUs are fully seated.
4. From the back of the cabinet, push on all the power plugs running from the rack-mounted devices to the PDUs to verify that the cables are fully seated.
5. Connect power to the PDUs. Use the NEMA L6-20, 280VAC, single-phase power cable.
   a. Plug the power cable into the PDU.
   b. Draw the power cord through the opening at the base of the cabinet.

© Copyright IBM Corp. 2003

   c. Plug the power cable into the wall outlet or other appropriate receptacle.
   d. Turn on the breaker switch for the customer's source power.
   e. Ensure the PDU circuit breakers are turned on.
6. Verify that all internal PDUs are powered up by viewing the power-on LEDs on components that are connected to the PDUs.
   - Servers display a blinking green light on the front panel when power is applied.
   - The following devices have no power switch and will power up automatically when the PDUs are powered up. Verify that these components have power applied.
     – KVM Switch
     – Cisco 10/100 Switch
     – Cisco Gigabit Switch
     – In-Reach Port Server

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

Power-on the Expansion cabinets

1. Power-on the Cluster Nodes manually with the power switch.
2. Verify that every RSA card has power. A green LED on the card will light up when the card has power.
3. Once an Expansion Cabinet is powered on, verify all front panel LEDs on the Cluster Nodes are steady ON, otherwise you will not see all the nodes in the configuration.

Repeat the procedure for every Expansion cabinet in the cluster before moving on to the Primary cabinet.

Power-on the Primary cabinet

Power-on the following devices in the order listed.

1. Storage Expansion Units
   - Flip the circuit breakers on the back of the device to the ON position.
2. Storage Controllers
   - Flip the circuit breakers on the back of the device to the ON position.
3. Management Node
   - Turn the power switch on the front of the device to the ON position.

     Verify that the system passes POST with no errors. During the boot process verify that the PXE BOOT attempts to run. If not, press F1 to enter the SETUP utility to add **Network** as a third boot option.

     Once the system boot has completed the screen will show the *No operating system* icon. If there are any yellow warning lamps on the Management Node, fix the underlying condition before continuing.
4. Cluster Nodes
   - Power-on Cluster Node 1 and verify that the system passes POST with no errors.

     During the boot process verify that the PXE BOOT attempts to run. If not, press F1 to enter the SETUP utility to add **Network** as a third boot option.

     Once the system boot has completed the monitor screen will show the *No operating system* icon. If there are any yellow warning lamps on the Cluster Node, fix the underlying condition before continuing.

     Repeat the procedure for all Cluster Nodes in the cluster.
5. Storage Nodes
   - Turn the power switch on the front of the device to the ON position.

- In order for the Storage Nodes to see them, the peripheral devices must be powered up and online before you power on the Storage Nodes.

6. Verify that all Cat-5 and fibre connections have a green link light.
7. Verify that every RSA card has power applied. A green LED on the card faceplate will light up when power is applied to the card. Connect your laptop 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**. Sign in to each RSA card using the web browser and verify that each is present. **Ping** each communication device (Cisco switches, In-Reach port server, power management module, and KVM switch).

If the system appears to be functionally sound, IBM will turn control over to the party installing the software.

Lights out or brown out

The following sequence should occur in a lights out scenario.

1. The system is up and running typical applications.
2. A lights out/brown out event occurs. The system powers down then powers up via an external source.
3. All nodes power back up to the last known state (on/off). If the last known state is on, then the nodes will boot to a login prompt.
4. Log files will show system restart events on nodes and on RSA cards. Check the following log files.
   - */var/log/messages*
   - */var/log/csm/installnode.log*
   - RSA event log
   - BIOS event log

Related topics

- Chapter 20, “Power Management Module replacement and configuration”, on page 111
- Chapter 9, “Cluster power down”, on page 55
- Appendix B, “Error Logs”, on page 125

Chapter 6. Software installation

Contents

These installation steps are used by IBM Global Services or the customer’s agent for the initial software set-up of a Cluster 1350. There are five basic steps to complete:

1. Install a supported distribution of Linux:
   - Red Hat Linux 7.3 or 8.0
   - Red Hat Linux Advanced Server 2.1
   - SuSE Linux 8.0 and 8.1
   - SuSE Linux Enterprise Server (SLES) 7 or 8
2. Install Cluster Systems Management (CSM)
3. Configure the storage nodes
4. Copy the system image to all nodes in the cluster
5. Test the configuration

Installation time is about 8 hours per cabinet.

Before you begin the software installation process, refer to “Software Version Matrix” to verify that you have all the required material.

ATTENTION!

The Cluster 1350 should be maintained only by system administrators experienced with Red Hat Linux, DHCP, NFS, and Linux networking and administration.

Software Version Matrix

The Cluster 1350 requires certain levels of a supported Linux distribution and Cluster System Management (CSM) in order to operate correctly. Before you begin the software installation process, make sure you have collected all the appropriate levels of operating system kernel, management software, device drivers, and other firmware needed for building a working system image. Table 10 shows the supported software and firmware versions.

Changing the versions of any components will adversely affect IBM’s ability to service and support the cluster.

Table 10. Cluster 1350 supported software and firmware versions - June 2003

Product	**Versions**
Red Hat Linux **-or-**	v7.3 or v8
Red Hat Linux Advanced Server**-or-**	v2.1
SuSE Linux **-or-**	v8 or v8.1
SuSE Linux Enterprise Server (SLES)	v7 or v8
Cluster System Management (CSM)	csm 1.3.1
Cisco 3508 GB Switch	12.1(9)EA1c, 3500L-C3H25-MZ-120-5.2-xv.1.bin

© Copyright IBM Corp. 2003

Table 10. Cluster 1350 supported software and firmware versions - June 2003 (continued)

Product	**Versions**
Cisco 35xx Ethernet Switch	12.1(9), 3500L-C3H25-MZ-120-5.3-wc.1.bin
Cisco 400x Ethernet Switch	12.1(13)EW1
Broadcom gigabit ethernet	bcm5700-ver 6.0.2
Ethernet Intel Single Port 10/100 Adapter	e1000-4.3.17.tar.gz
Fiber FAStT Adapter	BIOS v 3.01.31; Driver v6.01.00-fo
FAStT 700Raid Controller	ver 5.4
APC Power Switch	SNMP AOS v3.03, Masterswitch App v2.20
IBM NetBAY console switch	S2.0.0.I, F1.1.0, H0.0.10.02.00, A0.0.12
Myrinet 2000 Fiber PCI-2B card	GM-1.6.4
x335 BIOS ISMP Diagnostics Remote Supervisor Adapter F/W	v1.05 v1.03 v1.00 v1.03
BladeCenter BIOS ISMP Diagnostics	v1.03 15 v1.00
x345 BIOS ISMP Diagnostics Remote Supervisor Adapter F/W	v1.08 v1.04 v1.04 v1.17
x360 BIOS Diagnostics Remote Supervisor Adapter F/W	v1.08 v3.01 v1.08
ServeRAID Driver	6.10.20

Download link for drivers and firmware

If needed, use the following link to download the required drivers and firmware:

http://publib.boulder.ibm.com/cluster/1350downloads.htm

Install a supported distribution of Linux

The Cluster 1350 is shipped without an operating system. The customer or the customer's agent is responsible for securing a valid copy of the operating system for installation. Currently Red Hat Linux 7.3 or 8, Red Hat Linux Advanced Server 2.1, SuSE Linux 8 or 8.1, or SuSE Linux Enterprise Server (SLES) 7 or 8 are the only operating systems supported on the Cluster 1350.

Detailed installation instructions for each Linux distribution are provided by the manufacturer.

Installation of the operating system begins with the Management Node in the Primary cabinet.

Installing Red Hat Linux 7.3 or 8

The following steps provide a general path through the installation process and assume you will install Red Hat Linux 7.3 from the product CDs.

This process assumes you have successfully cabled together all devices within each cabinet and that you have correctly cabled together all cabinets. The software installation process begins with the Management Node in the Primary cabinet. Make sure the Management Node and all the switches are powered up. Cluster Nodes will be powered up later in the install process.

1. Power up the Management Node and press the F1 key. Make any necessary changes to set the boot order as follows:
   - Floppy diskette
   - CD-ROM
   - Fixed disk
   - Network
2. Ensure that hard disk 1 and 2 in the Management Node are set up as a RAID 1 mirror.

   To set up the disks as a RAID 1 mirror, reboot the Management Node and at the LSI Logic prompt press the **CNTRL + C** keys. Highlight the first drive listed with **Boot Order 0** and then highlight **Mirroring Properties**. Set drive 0 to primary and drive 1 to secondary. Press **ESC** and then **Enter** and save your configuration changes. Exit the configuration utility. You will get a message that the drives are resynching.
3. Insert the Red Hat 7.3 Linux product CD and restart the Management Node.
4. When prompted about the installation method you want to use select **CD-ROM**.
5. When prompted to choose an installation type, select **Custom**.

   You will use disk druid to manually partition the drive. First, delete all partitions on **sda**. Unhighlight all the drives listed by disk druid except for **sda** and create the following partitions:
   - A **/boot** partition with a size of 50 MB.
   - A **/csminstall** partition with 2 GB for each distribution of Red Hat Linux installed on the managed nodes.
   - A **/opt** partition of 2 GB.
   - A **/var** partition of 1 GB for each 128 nodes in the cluster. If there are fewer than 128 nodes in the cluster you should still set aside 1 GB. If there are more than 128 nodes you should increase the size of the partition to 2 GB or greater.
   - A **/home** partition of 2 GB, as a minimum. The customer should customize this to meet their expected needs.
   - **swap** defined as 2 GB.
   - A root partition mounted as **/** with an initial size of 4 GB. Make sure you select the **Fill to maximum allowable size** check box.
6. When prompted about installing a Boot loader, take the default selection and install the GRUB Boot loader in the Master Boot Record (MBR).

7. When prompted for a Boot loader password, enter one if security practices at the installation site require it. Make sure you record the password in a secure place.
8. When the Network Configuration screen appears enter the IP addresses for each device contained in the Management Node using the table shown below:

Network Name	**IP Range**	**Netmask**	**Network**	**Broadcast**	**Hostname**
eth0					
Management Node	172.20.0.1	255.255.0.0	172.20.0.0	172.20.255.255	mgt.cluster.com
eth1					
Management Node	172.30.0.1	255.255.0.0	172.30.0.0	172.30.255.25	mgt1.cluster.com
eth2 (customer network)					
Management Node	??	??	??	??	??

   If the Management Node in the Cluster 1350 is an x345, then **eth1** and **eth2**will not be available during the install process. Any network configuration for **eth1** and **eth2** must be done after the initial installation.
9. When prompted to install a firewall select **No**. Use a firewall for a public network interface only. The cluster and management networks must be unrestricted.
10. Set the default language and time zone for use on your system.
11. Set the root account password.
12. Create an **admin** account to log into after installation is complete.
13. When prompted to select the packages you want installed during the installation process check both the **Everything** and **Select individual packages** check boxes.
14. When the **Select individual packages** screen appears, select **System Environment** then **kernel** then **kernel-smp**. Make sure you deselect **kernel-pcmcia-cs**.

Once you have selected the packages you want installed the installation process will begin. When prompted for **Boot disk creation** insert a blank, formatted diskette into the drive and click **Next**. Label the diskette and keep it in a safe place.

Once the installation process has completed, you'll need to update the Intel network drivers.

Update the Management Node device drivers

In order for Red Hat Linux 7.3 to properly work on the Management Node and communicate with all the other devices in the cluster you need to download updates to the LSI drivers and Intel network drivers.

Go to http://publib.boulder.ibm.com/cluster/1350downloads.htm to download device drivers.

Download the appropriate files. Read and follow the installation instructions.

Update the configuration files

Now that the Management Node can communicate with other devices in the cluster take the following steps:

1. Update **/etc/modules.conf** with the following lines:
   ```
   alias eth0 e1000
   alias eth1 e1000
   alias eth2 e1000
   ```
2. Edit or create the **/etc/sysconfig/network-scripts/ifcfg-eth1** file and make sure it contains the following directives:
   ```
   DEVICE=eth1
   BOOTPROTO=static
   IPADDR=172.30.0.1
   NETMASK=255.255.0.0
   ONBOOT=yes
   ```
3. Edit or create the **/etc/sysconfig/network-scripts/ifcfg-eth2** file and customize it with site specific network information. It should contain the same directives as **ifcfg-eth1**.
4. Edit the **/etc/sysconfig/network** file. As a minimum, it should contain the following directives:
   ```
   NETWORKING=yes
   HOSTNAME=mgtnode.cluster.com
   GATEWAY=(site specific gateway IP--default route)
   GATEWAYDEV=eth2
   ```
5. Update the **/etc/resolv.conf** file with the following lines:
   ```
   search clusters.com (any other site specific domains)
   nameserver 172.20.0.1
   ```
6. Populate the **/var/named** directory using the **named** files found on the diskette shipped with the Cluster 1350.
7. Copy over the downloaded **/etc/named.conf** and */etc/hosts* files.
8. Edit the **/etc/named.conf** file and add any site specific nameservers to the forwarders section. For example:
   ```
   options {
           directory "/var/named';
           forwarders {
                   192.168.119.10;
                   192.168.119.11;
           };
           forward only;
   ```
9. Edit the **/etc/sysconfig/dhcpd** file. Change the DHCPDARGS directive to the following:
   ```
   DHCPDARGS="eth0 eth1"
   ```
10. Enable the named, dhcpd, and NFS services by issuing the following commands:
    ```
    # chkconfig named on
    # chkconfig dhcpd on
    # chkconfig nfs on
    ```
11. Remove kudzu by issuing the following command:
    ```
    # chkconfig – – del kudzu
    ```
12. Reboot the Management Node by issuing the following command:
    ```
    shutdown -r now
    ```

Installing SuSE Linux 8 or 8.1

The following steps provide a general path through the installation process and assume you will install SuSE Linux 8 from the product CDs.

This process assumes you have successfully cabled together all devices within each cabinet and that you have correctly cabled together all cabinets. The software installation process begins with the Management Node in the Primary cabinet. Make sure the Management Node and all the switches are powered up. Cluster Nodes will be powered up later in the install process.

1. Power up the Management Node and press the F1 key. Make any necessary changes to set the boot order as follows:
   - Floppy diskette
   - CD-ROM
   - Fixed disk
   - Network
2. Ensure that hard disk 1 and 2 in the Management Node are set up as a RAID 1 mirror.

   To set up the disks as a RAID 1 mirror, reboot the Management Node and at the LSI Logic prompt press the **CNTRL + C** keys. Highlight the first drive listed with **Boot Order 0** and then highlight **Mirroring Properties**. Set drive 0 to primary and drive 1 to secondary. Press **ESC** and then **Enter** and save your configuration changes. Exit the configuration utility. You will get a message that the drives are resynching.
3. Insert the SuSE Linux 8 product CD and restart the Management Node.
4. When prompted, select the **Language** you want to use for this installation.
5. Select **New Installation** and click **OK**.
6. Select **Partitioning**, then **Discard**, and then **Custom Partitioning**.

   You will manually partition the drive. First, delete all the partitions on **sda**. Create the following partitions:
   - A **/boot** partition with a size of 50 MB.
   - A **/csminstall** partition with 2 GB for each distribution of Red Hat Linux installed on the managed nodes.
   - A **/opt** partition of 2 GB.
   - A **/var** partition of 1 GB for each 128 nodes in the cluster. If there are fewer than 128 nodes in the cluster you should still set aside 1 GB. If there are more than 128 nodes you should increase the size of the partition to 2 GB or greater.
   - A **/home** partition of 2 GB, as a minimum. The customer should customize this to meet their expected needs.
   - **swap** defined as 2 GB.
   - A root partition mounted as **/** with an initial size of 4 GB.
7. Select **Software** and choose **Default system**. Then select **Detailed** and make sure **Network/Server** is selected. Select **Timezone** and change if necessary.

The operating system will start installing. At some point you will be asked to reboot the system to complete the installation.

After installation completes you can configure the Network interfaces. Enter the IP addresses for each device contained in the Management Node using the table shown below: Both **eth1** and **eth2**will not be available during the install process. Any network configuration for **eth1** and **eth2** must be done after the initial installation.

Network Name	IP Range	Netmask	Network	Broadcast	Hostname
eth0					
Management Node	172.20.0.1	255.255.0.0	172.20.0.0	172.20.255.255	mgt.cluster.com
eth1					
Management Node	172.30.0.1	255.255.0.0	172.30.0.0	172.30.255.25	mgt1.cluster.com
eth2 (customer network)					
Management Node	??	??	??	??	??

Update the configuration files

Now that the Management Node can communicate with other devices in the cluster take the following steps:

1. Update **/etc/modules.conf** with the following lines:

```
alias eth0 e1000
alias eth1 e1000
alias eth2 e1000
```

2. Edit or create the **/etc/sysconfig/network/ifcfg-ethx** file and make sure it contains the following directives:

```
BOOTPROTO="static"
BROADCAST="172.20.255.255"
IPADDR="172.20.0.1"
NETMASK="255.255.0.0"
NETWORK="172.20.0.0"
STARTMODE="onboot"
```

3. Edit or create the **/etc/sysconfig/network/ifcfg-eth2** file and customize it with site specific network information. It should contain the same directives as **ifcfg-eth1**.
4. Update the **/etc/resolv.conf** file with the following lines:

```
search clusters.com (any other site specific domains)
nameserver 172.20.0.1
```

5. Populate the **/var/named** directory using the **named** files found on the diskette shipped with the Cluster 1350.
6. Copy over the downloaded **/etc/named.conf** and **/etc/hosts** files.
7. Edit the **/etc/named.conf** file and add any site specific nameservers to the forwarders section. For example:

```
options {
        directory "/var/named';
        forwarders {
                192.168.119.10;
                192.168.119.11;
        };
        forward only;
```

8. Edit the **/etc/sysconfig/dhcpd** file. Change the DHCPD_INTERFACE directive to the following:

```
DHCPD_INTERFACE="eth0 eth1"
```

9. Enable the **named**, **dhcpd**, and **NFS** services by issuing the following command:

```
# chkconfig nfsserver on
# chkconfig named on
# chkconfig dhcpd on
```

10. Reboot the Management Node by issuing the following commands:

    ```
    shutdown -r now
    ```

Installing SuSE Linux Enterprise Server (SLES) 7 or 8

The following steps provide a general path through the installation process and assume you will install SuSE Linux Enterprise Server (SLES) 7 or 8 from the product CDs.

This process assumes you have successfully cabled together all devices within each cabinet and that you have correctly cabled together all cabinets. The software installation process begins with the Management Node in the Primary cabinet. Make sure the Management Node and all the switches are powered up. Cluster Nodes will be powered up later in the install process.

1. Power up the Management Node and press the F1 key. Make any necessary changes to set the boot order as follows:
   - Floppy diskette
   - CD-ROM
   - Fixed disk
   - Network
2. Ensure that hard disk 1 and 2 in the Management Node are set up as a RAID 1 mirror.

   To set up the disks as a RAID 1 mirror, reboot the Management Node and at the LSI Logic prompt press the **CNTRL + C** keys. Highlight the first drive listed with **Boot Order 0** and then highlight **Mirroring Properties**. Set drive 0 to primary and drive 1 to secondary. Press **ESC** and then **Enter** and save your configuration changes. Exit the configuration utility. You will get a message that the drives are resynching.
3. Insert the SuSE Linux Enterprise Server 7.2 product CD and restart the Management Node. At the installation prompt press the left **ALT** key and insert the LSI SLES update diskette.
4. When prompted, select the **Language** and the **Timezone** you want to use for this installation.
5. Select **New Installation** and click **OK**.
6. When prompted for **Preparing Hard Disk** select **Custom partitioning for experts**.

   You will manually partition the drive. First, delete all partitions on **sda**. Create the following partitions using the **ext2** file system.:
   - A **/boot** partition with a size of 50 MB.
   - A **/csminstall** partition with 2 GB for each distribution of Red Hat Linux installed on the managed nodes.
   - A **/opt** partition of 2 GB.
   - A **/var** partition of 1 GB for each 128 nodes in the cluster. If there are fewer than 128 nodes in the cluster you should still set aside 1 GB. If there are more than 128 nodes you should increase the size of the partition to 2 GB or greater.
   - A **/home** partition of 2 GB, as a minimum. The customer should customize this to meet their expected needs.
   - **swap** defined as 2 GB.
   - A root partition mounted as **/** with an initial size of 4 GB.

7. Select **Software Selection** and choose **Default system**. Then select **Detailed Selection** and make sure **Network/Server** is selected. Under **System Boot Configuration** select **Next**.
8. Reboot the node.
9. You can now configure the network using the values shown in the following table:

Network Name	IP Range	Netmask	Network	Broadcast	Hostname
eth0					
Management Node	172.20.0.1	255.255.0.0	172.20.0.0	172.20.255.255	mgt.cluster.com
eth1					
Management Node	172.30.0.1	255.255.0.0	172.30.0.0	172.30.255.25	mgt1.cluster.com
eth2 (customer network)					
Management Node	??	??	??	??	??

After you have configured the network, network initialization will fail. To correct this problem, update the Intel e1000 drivers and update SLES 7 or 8 to the latest level kernel available.

Update the configuration files

Now that the Management Node can communicate with other devices in the cluster take the following steps:

1. Update **/etc/modules.conf** with the following lines:

```
alias eth0 e1000
alias eth1 e1000
alias eth2 e1000
```

2. Edit or create the **/etc/rc.config** file and make sure it contains the following directives:

```
#
# IP Addresses
#
IPADDR_0="172.20.0.1"
IPADDR_1="172.30.0.1"
IPADDR_2=""
IPADDR_3="

#
# Network device names (e.g. "eth0")
#
NETDEV_0="eth0"
NETDEV_1="eth1"
NETDEV_2=""
NETDEV_3=""


#
# Parameters for ifconfig, simply enter "bootp" or "dhcpclient" to use
# the respective service for configuration.  Sample entry for ethernet:
# IFCONFIG_0="192.168.81.38 broadcast 192.168.81.63 netmask 255.255.255.224"
#
IFCONFIG_0="172.20.0.1 broadcast 172.20.255.255 netmask 255.255.0.0"
IFCONFIG_1="172.30.0.1 broadcast 172.30.255.255 netmask 255.255.0.0"
IFCONFIG_2=""
```

```
IFCONFIG_3=""


#
#Hostname of the system (full name)
# If zero, and bootp is used above, bootp will also set the hostname
# (e.g. "riemann.suse.de" or "hugo.linux.de")
# Don't forget to edit your etc/hosts file appropriately.
#
FQHOSTNAME="mgtnode.cluster.com"'
```

3. Update the **/etc/resolv.conf** file with the following lines:

   ```
   search cluster.com (any other site specific domains)
   nameserver 172.20.0.1
   ```

4. Populate the **/var/named** directory using the **named** files found on the diskette shipped with the Cluster 1350.
5. Copy over the downloaded **/etc/named.conf** and **/etc/hosts** files.
6. Edit the **/etc/named.conf** file and add any site specific nameservers to the forwarders section. For example:

   ```
   options {
           directory "/var/named';
           forwarders {
                   192.168.119.10;
                   192.168.119.11;
           };
           forward only;
   ```

7. Edit the **/etc/rc.config.d/dhcpd.rc.config** file. Change the DHCPD_INTERFACE directive to the following:

   ```
   DHCPD_INTERFACE="eth0 eth1"
   ```

8. Enable the **named**, **dhcpd**, and **NFS** services by issuing the following command:

   ```
   # chkconfig nfsserver on
   # chkconfig named on
   # chkconfig dhcpd on
   ```

9. Reboot the Management Node by issuing the following commands:

   ```
   shutdown -r now
   ```

Install Cluster Systems Management (CSM)

Once you have installed Linux on the Management Node you should install CSM. This section covers:

- CSM pre-installation tasks
- Installing CSM on the management node

Pre-installation tasks

1. Download the CSM 1.3 Software Planning and Installation Guide from:

   http://www-1.ibm.com/servers/eserver/clusters/library/csmsetup.html

   The CSM guide has detailed instructions that you will need for installing CSM on the Management Node.
2. Prior to installing CSM, ensure you have located the Cluster 1350 Install Data Diskette that was shipped with the cluster. The diskette contains node definition information needed by CSM.

   The Cluster 1350 Install Data Diskette holds the node definition information for up to 64 cluster nodes. An extra diskette is required for each additional group or partial group of 64 cluster nodes. All diskettes are consecutively numbered so you can put them in the proper order before continuing with the next step.

Two complete sets of diskettes are included with each cluster. If any diskettes are missing or damaged, contact IBM Support for information on how to proceed.

3. Each diskette contains an **xcat tab** file that must be converted for use with CSM. If the cluster has 64 cluster nodes or fewer you will have only one **xcat tab** file to convert . Go to "Converting xcat tab file for use with CSM" and continue with the steps shown there.

   If the cluster has more than 64 cluster nodes you will have more than one diskette containing an **xcat tab** file. The information contained on the diskettes must be combined into one large file before converting the file for use by CSM. Continue with the next step.
4. Make sure the first 1350 Install Data Diskette is in the diskette drive of the Management Node and issue the following commands:

   ```
   mount/dev/fd0/
   mnt/floppy
   cp */mnt/floppy/ /tmp
   ```

   Continue by copying the **xcat tab** file from each 1350 Install Data Diskette to the */tmp* directory. When all the files have been copied, continue with the next step.
5. Download the **merge-ic.sh** script from: http://publib.boulder.ibm.com/cluster/1350downloads.htm and place the script in the */tmp* directory. Issue the following command:

   ```
   merge-ic.sh tarfile-1 tarfile-2 tarfile-3...
   ```

   where *tarfile-1 tarfile-2 tarfile-3...* are the actual names of the tarfiles copied to */tmp* from the 1350 Install Data Diskettes. After the **merge-ic.sh** has finished go to "Converting xcat tab file for use with CSM" and continue with the steps shown there.

Converting xcat tab file for use with CSM

To convert the **xcat tab** file for use with CSM take the following steps:

1. For clusters with 64 nodes or fewer insert the 1350 Install Data Diskette in the diskette drive of the Management Node and issue the following commands:

   ```
   mount/dev/fd0
   tar xvzf/mnt/floppy/ic.tgz -C/opt/xcat
   umount/dev/fd0
   ```

   Remove the diskette from the drive and store the diskette in a safe place.

   For clusters with more than 64 nodes you must use the merged file you created that is located in the */tmp* directory. Issue the following command:

   ```
   tar xvzf/tmp/ic.tgz -C/opt/xcat
   ```
2. Install the Enhanced Cluster Tools (ECT) for Linux Cluster Systems Management. If you don't already have a copy, download one from http://www.alphaworks.ibm.com/tech/ect4linux
3. Issue the following commands:

   ```
   rpm -ivh csm.ect-1.3.1.0-4.i386.rpm
   XCATROOT=/opt/xcat/opt/xcat/sbin/xcat2csm -F /tmp/nodedef
   rm -rf /opt/xcat
   ```

Installing CSM on the Management Node

Perform the following to install CSM on the **Management Node**:

1. Set the following environment variables:
   - Make sure **/opt/csm/bin** is listed in the PATH environment variable
   - Make sure **/opt/csm/man** is listed in the MANPATH environment variable
   - Set the environment variable **RCONSOLE_FONT** to the font of your choice.
   - Set the environment variable **DSH_REMOTE_CMD** to the remote shell of choice (default=**ssh**)

   You can put these environment variables into the **/etc/profile.d/csm.{sh,csh}** shell scripts so they are automatically loaded.
2. Refer to the CSM Software Planning and Installation Guide for detailed installation instructions. The guide is located at: http://www-1.ibm.com/servers/eserver/clusters/library/csmsetup.html
3. Once CSM is installed, take the 1350 node definition information contained in the **nodedef.install** file into **/opt/csm/install**.
4. Issue the command **definenode –f nodedef.install**.
5. Issue the command **lsnode –l** and check to see that all the nodes you were expecting show up in the listing.
6. Update CSM with the new required drivers. The required drivers are located at:http://techsupport.services.ibm.com/server/cluster2/fixes/csmdriverdownload.html.
7. Issue the command **rpower –a query** and check that a valid status (either **on** or **off**) is returned for each node defined.
8. Issue the command **rconsole –a** to check that you have valid CSM console definitions and connectivity to the terminal server.
9. Modify the appropriate **kscfg.tmpl.*** file in **/opt/csm/install** before running the kickstart generation script.

   Create the CSM kickstart.

   For Red Hat 7.3 issue the following command: **csmsetupks**

   For SuSE Linux 8 issue the following command: **csmsetupsis —a —p /cdrom:/tmp/sis —A suse**

   For SLES 7.2 issue the following command: **csmsetupsis —a —p /cdrom:/tmp/sis —A sles**
10. Follow the prompts and load the Linux CD-ROMs as required.

ATTENTION!

Once **csmsetupks** or **csmsetupsis** has successfully run for each of the nodes you are ready to install the operating system on them.

Before you continue, either power off the FAStT controllers or disconnect the fibre cable that runs to each FAStT controller.

Configure the Storage nodes

Prerequisites

The procedure assumes the following prerequisites:

- Red Hat 7.3 is installed and running from a local drive.
- The FAStT Storage Server is properly configured and connected to a Host Bus Adapter (HBA).
- The FAStT drives in the Storage Server are configured into different RAID groups, Storage Groups, and LUNS via the software provided with the FASTt Storage Server.

Issues

Because of the way the Red Hat 7.3 loads SCSI drivers and assigns them to **/dev/sda**, **/dev/sdb**, and so on. Problems can result if more than one SCSI host adapter board (Adaptec SCSI controller for local drives and Qlogic HBA for Triton connection) is installed on the system and you use the **scsi_hostadapter** alias. When the system is rebooted, the operating system will detect the Qlogic controller prior to detecting the Adaptec, which will cause the system to panic. To avoid this problem, follow the "Installation procedure" and modify the order of the contents of the **/etc/modules.conf** file.

Installation procedure

1. **If not already done, power down the storage controllers or disconnect the fibre cable running to each storage controller.**
2. For Red Hat 7.3 or 8.0 edit the **/etc/modules.conf** file to put the host adapters in the correct order and to add the parameter `scsi_mod max_scsi_luns` to the file.

ATTENTION!

Because the system is running a modular kernel, the Adaptec SCSI device driver ("alias scsi_hostadapter aic7xxx") must be probed before any other SCSI adapters so that the kernel will be able to find the root device during the initialization phase. Also, the Qlogic Qla driver ("alias scsi_hostadapter2 qla2x00") must be the last SCSI host adapter listed in the file.

If there are other SCSI host adapter boards installed on your system and the scsi_hostadapter alias is used, define a different alias for the qlogic Qla driver and make sure to add it after the other SCSI modules so this doesn't cause the SCSI devices names already in use to be renumbered on next boot. You can do this by appending a number at the end of the scsi_hostadpter word, for example, **alias scsi_hostadapterN qla2x00** [where N is a number from 1-9].

For example:

```
Original modules.conf:
alias eth0 e1000
alias scsi_hostadapter qla2x00
alias scsi_hostadapter1 aic7xxx
alias scsi_hostadapter2 ips
alias eth1 e1000
alias parport_lowlevel parport_pc
alias scsi_hostadapter3 aic7xxx
alias scsi_hostadapter4 aic7xxx
alias usb-controller usb-ohci

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

For SuSE Linux 8 or 8.1 and SLES 7.2 or 8 edit the **/etc/modules.conf** file to ensure it contains the following lines:

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

For SuSE Linux 8 or 8.1 edit the **/etc/sysconfig/kernel** file to contain the following line:

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

For SLES 7.2 or 8 edit the **/etc/rc.config** file to contain the following line:

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

3. For Red Hat 7.3 or 8.0 rebuild the two **initrd** images (**mkinitrd** will not allow you to make a ramdisk image if it detects one already present with the same name, so the first two commands will rename the old images):

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

   For SuSE 8 or 8.1 and SLES 7.2 or 8 run the **mkinitrd** command to create a **boot/initrd** directory and then run **lilo**.
4. If an RSA adapter was installed, reboot and load the setup floppy or CD-ROM to configure the network. Assign the same configuration information for the RSA adapter (name, IP, hostname) as used before.

ATTENTION!

Refer to this site to download RSA and ASM Process or Firmware Update Diskette utility:
http://www.pc.ibm.com/qtechinfo/MIGR-4JTS2T.html

5. Configure the kernel (if you have custom modifications).
6. Reboot the node.

Copy the system image out to all nodes in the cluster

Because of the way the Red Hat 7.3 and 8 loads SCSI drivers and assigns them to **/dev/sda**, **/dev/sdb**, and so on, problems can result if more than one SCSI host adapter board (Adaptec or LSI SCSI controller for local drives and Qlogic HBA for Triton connection) is installed on the system. The QLogic HBA will typically be seen first by the install process. Follow the “Installation procedure” on page 47 and modify the order of the contents of the **/etc/modules.conf** file.

Attempting to copy the system image out to the nodes while a FAStT controller is still powered up and connected may cause data corruption on the first logical disk device in FAStT subsystem. Ensure the FAStT controllers are powered down or all fibre cables for the FAStT controllers are disconnected from the back of each controller before starting the install process.

To copy the system image out to all nodes in the cluster, take the following steps:

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

   ```
   rconsole -n {node list}
   ```
2. Run the **installnode** command for each node being installed:

   ```
   installnode {node list}
   ```

Once the operating system is installed on the storage nodes, reconnect the fibre cable to the FAStT controllers. Reboot the storage nodes to see any configured LUNs.

Test the configuration

1. Boot and log on to the Management Node as user **root**.

2. Log on to the Storage Nnodes and verify disk configuration. This can be done by using the **fdisk -l** command.
3. If present, configure the modem according to the manufacturer's instructions.
4. The system is now ready for the customer to connect their network cables.

Chapter 7. Cluster management

Contents

IBM Cluster Systems Management provides a powerful way to administer the daily operations of a Cluster 1350.

For more information on such topics as Overview, Monitoring, Remote Control, Set-up, and Technical Reference, refer to:
http://www.ibm.com/servers/eserver/clusters/library/linux.html

© Copyright IBM Corp. 2003

Chapter 8. Remote access

Contents

Remote power

The command **rpower** boots and resets hardware, powers hardware on and off, and queries node power state. The syntax is:

```
rpower [-a] [-h] [-n host[,host...]] [-N Node_group[,Node_group...]]
 [-v] on | off | reboot | query | resetsp
```

Remote console

The Remote console function is provided using the serial ports of the x335 servers and terminal servers. This provides remote access to nodes before the operating system is installed or when network access to the servers is not available or failed. The terminal servers are required by the Install function in CSM and must be included to enable the Remote console function.

Each rack in the configuration includes one or two terminal servers to connect each node in the rack via a DB9 to RH45 serial cable. The terminal servers are LAN connected to the Management VLAN.

The Remote console function is accessed via the **rconsole** command. This command opens a remote console for each node specified with the command. The syntax is:

```
rconsole [-a] [-h] [-n host[,host...]] [-N Node_group[,Node_group...]
```

© Copyright IBM Corp. 2003

Chapter 9. Cluster power down

Contents

Power down the system

Because the operating system is installed on each node, the power down procedures are relatively simple.

1. Log off Cluster Nodes and Storage Nodes. If CSM is installed on the Management Node, execute the **rpower off -A** command to power down the Cluster Nodes. If CSM is not installed, power down the Cluster Nodes manually with the power switch.
2. Power down the following devices in the order listed.
   a. Storage Nodes
   b. Management Node
   c. Storage Controllers
   d. Storage Expansion Units
3. Power down the PDUs or unplug, from the PDU, the devices that have no power switch.
   - To power down the PDUs, unplug them from the wall outlet.

   Note: The following devices have no power switch and must be unplugged if the PDUs are not powered down.
   - KVM Switch
   - Cisco 10/100 Switch
   - Cisco Gigabit Switch
   - iTouch terminal server
4. Unplug the PDUs from the wall outlet.

Lights out or brown out

The following sequence will occur in a lights out scenario:

1. The system is up and running typical applications.
2. A lights out/brown out event occurs. The system powers down then powers up via an external source.
3. All nodes power back up to the last known state (on/off). If the last known state is on, then the nodes will boot to a login prompt.
4. Log files will show system restart events on nodes and on RSA cards. Check the following log files.
   - */var/log/messages*
   - */var/log/csm/installnode.log*
   - RSA event log.
   - BIOS event log

Related topics

- Chapter 5, “Power up the cluster”, on page 31
- Appendix B, “Error Logs”, on page 125

© Copyright IBM Corp. 2003

Part 3. Service

© Copyright IBM Corp. 2003

Chapter 10. Hardware/software problem determination

Contents

How to use this information

This chapter helps diagnose problems associated with the eServer Cluster 1350. Cluster 1350 is an integrated Linux Cluster that includes IBM and Third Party hardware and software components like server nodes and associated service processors, storage and networking subsystems, plus Cluster Systems Management (CSM) and General Parallel File System (GPFS) software.

Problem determination involves identifying the likely cluster component where the problem might have occurred, and following the relevant problem determination steps for that component.

This chapter will aid in the diagnosis of problems down to the component level. Once a failing componenet is identified you should refer to the component's product documentation for further actions. Links to product web sites and online product documentation are provided in this chapter as appropriate.

Diagnosing hardware/software problems in a clustered environment requires a basic understanding of how the components of the eServer Cluster 1350 function together.

The cluster consists of:
- One or more 19″ racks.
- From 4 to 512 Cluster Nodes. The nodes of the cluster may be an x335 or BladeCenter containing at least four Blade servers. 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 (an x345) for cluster systems management and administration.
- A Management Ethernet VLAN used for secure traffic for hardware control.

  The Management Ethernet VLAN is used for management traffic only. It is logically isolated for security using the VLAN capability of the Cisco Ethernet switches, and is only accessible from the Management Node. The Cluster VLAN and Management VLANs share the same physical Cisco switches.
- A Cluster VLAN used for other management traffic and user traffic. Cisco switches integrated with the cluster are used for the Management Ethernet VLAN and the Cluster Ethernet VLAN.
- Service Processor networks. All nodes in the cluster are connected via daisy-chained service processors (x335) and/or Remote Supervisor Adapter cards. The first node in a daisy-chain must have a Remote Supervisor Adapter which is Ethernet connected to the Management Ethernet VLAN.
- A Terminal Server network for Remote Console, using the MRV In-Reach terminal server. Optionally, the customer may elect to include an additional network.
- A high-performance Myrinet 2000 cluster interconnect, or an additional 10/100 Ethernet.

© Copyright IBM Corp. 2003

- The customer may elect to configure a subset of Cluster Nodes with additional external storage - for example as Storage Nodes for GPFS. This may be a Fiber Channel solution (using a FAStT storage subsystem).
- A supported distribution of the Linux operating system.
- Cluster management software such as CSM.

CSM maintains a database of configuration information about the nodes that are configured in the Cluster 1350. The following CSM command run on the management server will display this node information: **lsnode -Al**

The output of this command provides a great deal of information about each node such as its type, model number, serial number, and hostname. The output also provides important information that relates each node to the terminal server network and service processor network. For the terminal server network, the output includes the console server hostname and the console port number to which the node is connected. For the service processor network, the output includes the hostname of the Falcon card to which the node is connected and the internal service processor name for the node.

CSM distinguishes between the management node, pre-managed nodes, and managed nodes. Pre-managed nodes are nodes that have been added to the configuration but are not yet ready to be managed, for example, because they have not yet been installed. Running **lsnode -P** on the management server will generate a list of pre-managed nodes in the configuration while **lsnode** only lists the managed nodes. Running **mgmtsvr** from any node will display the management server.

Isolating network, node, and Linux problems

Cluster 1350 nodes are connected over a 10/100 Mbit Ethernet Cluster network. A Cluster 1350 may 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.

Note the results and compare the symptoms seen to Table 11

Table 11. Troubleshooting the shared VLAN

Symptom	Action

Table 11. Troubleshooting the shared VLAN (continued)

1. Can **ping** the Storage Node from the Management Node but cannot **ping** the Cluster Nodes 2. Can **ping** the Cluster Nodes from the Management Node but cannot **ping** the Storage Nodes 3. Cannot **ping** either the Storage Nodes or the Cluster Nodes. 4. Cannot **ping** the Cluster Nodes in one of the expansion cabinets.	1. Verify links between the Management Node, Storage Nodes, Cisco 3550, 3500, and 400x switches. 2. Reboot the Management Node and press F1 to enter SETUP. Verify that Ethernet devices are turned on. 3. Verify correct driver level for 1GB fibre Ethernet. Verify status using the **ifconfig** command. 4. Verify the correct level of the driver is installed for the 1 GB fibre ethernet 5. Check the Cisco (3500,4003, or 4006) switch for green status lights or system and status LEDs. If the green lights are on the switch is OK. 6. Verify 1 GB fibre ethernet connections are good by swapping a known good cable to isolate the failing device. 7. Replace the failing fibre cable, GBIC, or network interface card.

If following the steps in Table 11 on page 60 did not correct the problem, continue with the steps shown in “Cluster with one network”.

Cluster with one network

Ping failure on one or some nodes: If one or more nodes experience a **ping** failure, it indicates a problem with the node hardware or software.

1. Attempt to **telnet** to the node via the serial console or KVM and verify the node is operational.
   a. If **telnet** succeeds, check the *syslog* for errors.
      1) If there are errors, go to“Isolating software problems” on page 69 and continue with the steps shown there.
      2) If there are no errors, it indicates a network problem. Go to Table 12 and continue with the steps shown there.
   b. If **telnet** fails, it indicates a node hardware problem. Go to “Isolating hardware problems” on page 63 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 toTable 12 and continue with the steps shown there.
- Network adapter on the Management Node.
- DHCP configuration
- Network configuration
- Cisco blade failure

Table 12. Network troubleshooting for a cluster with one network

Symptom	**Action**

Table 12. Network troubleshooting for a cluster with one network (continued)

1. Cannot **ping** a node or nodes on the cluster network from the Management Node, yet the **rconsole** command and access from the KVM work correctly.	1. Use the **ifconfig** command to verify that the IP settings are correct. 2. Verify that the cables are fully plugged into the switch and node, and that everything is plugged into the correct port. Refer to the cabling information printed on each cable label and "VLAN options" on page 17 if you are unsure where a cable belongs. Verify that the link lights are on. 3. Swap ports on the Ethernet switch with a Cluster Node port you know is working. 4. Verify the Ethernet switch port is configured for the Management VLAN.

Cluster with two networks

Ping failure on one or some nodes: If one or more nodes experience a **ping** failure, it indicates a problem with the node hardware or software.

1. Attempt to **telnet** to the node via the serial console or KVM and verify the node is operational.
   a. If **telnet** succeeds, check the *syslog* for errors.
      1) If there are errors, go to "Isolating software problems" on page 69 for software problem resolution.
      2) If there are no errors, it indicates a network problem. Go to Table 14 on page 63 and continue with the steps shown there.
   b. If **telnet** fails, it indicates a node hardware problem. Go to "Isolating hardware problems" on page 63 for problem resolution.

Ping failure on only one network: If **ping** failures occur on one network but not on the second network, it indicates a problem on the network adapter on the Management Node for the network where the failure occurred.

Ping failure on one or both networks:

1. Verify that all communication devices on the network are powered on and that each device has a green status light on both ends of the connection.
2. Verify with support that correct IP Address, Net Mask, and Gateway settings for each device that fails to function in the network.
3. Use the **ifconfig** command to determine the IP Address scheme of each node and compare it to the factory defaults shown in Table 13

Table 13.

Device	IP address	Host Name
Management Node	172.20.0.1 172.30.0.1	**eth0**–mgtnode.cluster.net **eth1**–mgtnode-eth1
Storage Node	172.20.1.1	storage001
First FAStT Storage	172.20.2.1	
Second FAStT Storage	172.20.2.2	
x335 Cluster Nodes	172.20.3.1	node001...node*xxx*
BladeCenter Ethernet Switch Module	172.20.90.1	

Table 13. (continued)

Myrinet switch	172.20.10.1	myri001
First iTouch port server	172.30.20.1	ts001
Second iTouch port server	172.30.20.2	ts002
RSA cards (bottom card)	172.30.30.1	rsa001
RSA cards (next card)	172.30.30.2	rsa002
RSA cards (Myrinet switch)	172.30.30.3	
Cisco 4003 switch (console management)	172.30.80.1	cisco4003–001
Cisco 10/100 switch	172.30.40.1	cisco3550–001
Cisco GB switch	172.30.50.1	cisco3508–001
APC Masterswitch	172.20.60.1	apc001
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.

1. Attempt to **telnet** to the node via the serial console.
   a. If **telnet** succeeds, check the *syslog* for errors.
      1) If there are errors, go to "Isolating software problems" on page 69 for software problem resolution.
      2) If there are no errors, it indicates a user application problem.
   b. If **telnet** fails, try using a serial communications program like Hyperterminal to connect to the node. If you still can not connect it indicates a node hardware problem. Go to "Isolating hardware problems" for problem resolution.

Table 14. Network troubleshooting for a cluster with two networks

Symptom	Action
1. Cannot **ping** a node or nodes on the cluster network from the Management Node, yet the**rconsole** command and access from the KVM work correctly.	1. Use the **ifconfig** command to verify that the IP settings are correct. 2. Verify that the cables are fully plugged into the switch and node, and that everything is plugged into the correct port. Refer to the cabling information printed on each cable label and"VLAN options" on page 17if you are unsure where a cable belongs. Verify that the link lights are on. 3. Swap ports on the Ethernet switch with a Cluster Node port you know is working. 4. Verify the Ethernet switch port is configured for the Management VLAN.

Isolating hardware problems

Node checks

Table 15. Troubleshooting the remote console network

Symptom	Action

Table 15. Troubleshooting the remote console network (continued)

1. Cannot execute any **rconsole** command to any Cluster Node. 2. Cannot execute any **rconsole** commands to get an active terminal session.	1. Verify 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. Verify that the serial port at the node is attached to the correct port on the terminal server by using the CSM command **LSNODE-aI <NodeName>**. Refer to the information listed under the ConsolePortNum category. 3. Follow steps 1 through 9 of “Configuration and setup after device replacement” on page 90, then at the ***IN-Reach_Priv>*** prompt type *show port <portnumber>* to verify the settings of all suspect ports against ports that are working correctly. 4. Verify the InReach terminal server is powered up and functional by pinging the unit at 172.30.20.1 5. Ensure the serial port (COM 1) is configured to redirect output to the terminal server. Take the following steps: • Restart the node and watch the monitor screen. • When the message **Press F1 for Configuration/Setup** appears, press **F1**. • From the main menu, select **Devices and I/O Ports** then press **Enter**. • Verify that Serial Port A is set to **Port 3F8, IRQ 4**. • Select **Serial Port A**. • Select **Remote Console Redirection**. • Verify 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. For the x335 and x345 only, run diagnostics against the serial port to certify function. 7. Swap out the cables and dongle with cables and dongle known to be good.

If the above procedures do not correct the problem you may have a problem with a port on the terminal server.

Try a different port and retest. Issue the CSM command **lsnode —AI <nodename> lgrep** and record the port information shown. Move the cable to a new port and change the port number using the CSM command **chnode <nodename> ConsolePortNum=***xx* where *xx* is the new port number. If the symptom persists, go to "Checking service processor logs" on page 72 and check the service processor log.

Hardware problem in service processor log: Go to "Node" on page 72 for node problem determination.

Amber LED lit on node: Service processor log may be full. The log is cleared by connecting to the service processor via the Falcon card. Otherwise, go to "Node" on page 72 for node problem determination.

rpower to node fails:

- Check service processor connection.
- Try the **rpower -a on** command again.
- Go to "Checking service processor logs" on page 72 and check the service processor log.
- Use the web interface or telnet to each RSA adapter on the system and then connect to the service processor on the Cluster Node individually through the ″Remote ASM Access″ menu.
- If the Cluster Node is not in the list, use the node's firmware diskette to diagnose the problem.

Service processor network

Table 16. Troubleshooting the service processor network

Symptom	Action

Table 16. Troubleshooting the service processor network (continued)

1. The **rpower —a query** command does not return with the status of all nodes 2. Cannot see all the nodes when managing remote ASMs. 3. Cannot connect to individual RSA cards using browser.	1. Check the physical connections on the RS485 network and check for errors.. 2. From the Management Node use the web browser and try to link up with the failing node through the failing node's RSA card. 3. Check that the RSA network is properly terminated. When more than one node is connected, terminators should be plugged into the empty port on the dongle and in the second RS485 port of the last node in the chain. 4. Swap the internodal Cat 5 cable on the unresponsive node with a known good cable. Also, replace the dongle if a problem is suspected. 5. Swap the KVM/RS485 cable (on the x335 only) with a known good cable. Also, replace the dongle if a problem is suspected. 6. Verify the RSA configurations/IP settings with support. 7. Verify the 10/100 ethernet link from the RSA card to the Cisco 3550 or 400x switch. 8. Flash the ASM to the latest firmware level. 9. Flash the RSA to the latest firmware level. 10. Check RSA configurations using the firmware update diskette.

If following the steps in Table 16 on page 65 did not correct the problem, continue with the steps shown in "RSA card connection failure".

RSA card connection failure:

1. Verify node has power using **rpower query** command.
   a. If node has power, ping the RSA card using the *HWControlPoint* field in **lsnode** output.
      1) If **ping** succeeds, reset adapter. If adapter connection continues to fail after it has been reset contact IBM support.
      2) If **ping** fails, check network connection.
   b. If node does not have power, check power connections.
2. If network connection LED is on for the RSA card at the Ethernet switch, go to "Resetting RSA cards" on page 72 and reset the RSA adapter.

Node connection or command failure: If the RSA card connection is working, but the node connection or commands issued to the node fail:

- Connect to the RSA card and verify node list.
- Verify cabling.
- Go to "Node checks" on page 63 and perform node checks.

Storage checks

Table 17. Troubleshooting the Fibre Storage network

Symptom	Action
1. Cannot see disk drives from the Storage Node	1. Reboot the Storage Node and press the **Alt/Q** keys to go into Qlogic setup. Verify that the 700 FastT is a listed device. 2. Check the Fibre connections between the server HBA and the hubs on the 700 FastT. Green connection lights hould be on. 3. Check cabling on the 1742 FastT outbound hubsw to the storage expansion units. Look for links lights and proper cabling. Also verify that all transfer rate speed switches are set to 2 gigabytes. 4. Check blade and ESM firmware levels and update to current levels.

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

File system failure: Check disks using **fdisk -l**:

- If **fdisk -l** completes, go to "GPFS" on page 72 and continue with the file system problem determination.
- If **fdisk -l** reports missing disks, check that the adapter device driver is configured
  - If the adapter device driver is configured, go to "Storage" on page 73 and continue with storage subsystem problem determination.
  - If the adapter device driver is not configured, check the adapter hardware and then go to "Linux" on page 72 and continue with Linux problem determination.

PFA alert indicates internal disk: Go to "Storage" on page 73 and perform disk problem determination.

I/O errors in syslog: Execute problem determination for the indicated disk, adapter or storage subsystem.

Ethernet checks

Ping failure over Ethernet: Check nodes using **rconsole** command or **ping** nodes via Myricom:

- If node responds, go to "Cisco switches" on page 73 and continue with Ethernet problem determination.
- If command fails, go to "Node" on page 72 and continue with Node checks.

Myricom checks

Ping failure over Myricom: Check nodes using **rconsole** command or **ping** nodes via Ethernet:

- If node responds, go to "Myrinet" on page 73 and continue with Myricom problem determination.
- If command fails, go to "Node" on page 72 and continue with Node checks.

Terminal server checks

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

Symptom	Action
1. Unable to execute **rconsole** commands and get an active terminal session.	1. Check conncection of cables and connectors at the nodes and the InReach terminal server. 2. Verify the InReach terminal server is powered up and functional by pinging the unit at 172.30.20.1. 3. Follow steps 1 through 9 of "Configuration and setup after device replacement" on page 90, then at the ***IN-Reach_Priv>*** prompt type *show port <portnumber>* to verify the settings of all suspect ports against ports that are working correctly. 4. Verify the InReach terminal server is powered up and functional by pinging the unit at 172.30.20.1 5. Ensure the serial port (COM 1) is configured to redirect output to the terminal server. Take the following steps: • Restart the node and watch the monitor screen. • When the message **Press F1 for Configuration/Setup** appears, press **F1**. • From the main menu, select **Devices and I/O Ports** then press **Enter**. • Verify that Serial Port A is set to **Port 3F8, IRQ 4**. • Select **Serial Port A**. • Select **Remote Console Redirection**. • Verify 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 certify function.

If following the steps in Table 18 did not correct the problem, continue with the steps shown in "Terminal server connection failure" on page 69.

Terminal server connection failure: Check nodes using **telnet** command or **ping** nodes via Ethernet:

- If commands fail, go to “Node” on page 72 and continue with Node checks.

KVM network

Table 19. Troubleshooting the KVM network

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

File system failure

Check disks using **fdisk -l**:

- If **fdisk -l** completes, go to “GPFS” on page 72 and continue with the file system problem determination.
- If **fdisk -l** reports missing disks, check that the adapter device driver is configured
  - If the adapter device driver is configured, go to “Storage” on page 73 and continue with storage subsystem problem determination.
  - If the adapter device driver is not configured, check the adapter hardware and then go to “Linux” on page 72 and continue with Linux problem determination.

PFA alert indicates internal disk

Go to “Storage” on page 73 and perform disk problem determination.

I/O errors in syslog

Execute problem determination for the indicated disk, adapter or storage subsystem.

Isolating software problems

Operating system checks

Node unresponsive: If the node does not respond to ping or the serial console and there are no relevant entries in syslog or hardware logs, go to “Linux” on page 72 to continue with the problem determination process.

Adapter device driver not configured: If the device driver is not configured, and there are no adapter hardware problems reported, go to "Linux" on page 72 to continue with the problem determination process.

CSM checks

Events not logged or actions not taken: Using the ERRM CLI, 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 for managed nodes which are currently up and reachable from the management server, follow problem determination section in Monitoring HOWTO located at: http://www-1.ibm.com/servers/eserver/clusters/library/csmadm.html

Differences in node lists: Output from 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. Likely configuration or network problem, follow problem determination section in Monitoring HOWTO located at: http://www-1.ibm.com/servers/eserver/clusters/library/csmadm.html

netstat output incomplete: The command **netstat -an | grep rmc** on the management server does not show *ESTABLISHED TCP* connections to each currently managed node which is up. Follow problem determination section in Monitoring HOWTO located at: http://www-1.ibm.com/servers/eserver/clusters/library/csmadm.html

RMC not running: The command **lssrc -ls ctrmc** shows that RMC is not running on the management server. Follow problem determination section in Monitoring HOWTO located at: http://www-1.ibm.com/servers/eserver/clusters/library/csmadm.html

lsrsrc reports errors: The command **lsrsrc -ab IBM.[Host|FileSystem]'** which checks that HostRM and FSRM will run on the management server reports errors. Follow problem determination section in Monitoring HOWTO located at: http://www-1.ibm.com/servers/eserver/clusters/library/csmadm.html

lsaudrec reports errors: The command **lsaudrec** which checks that AuditRM will run on the management server reports errors. Follow problem determination section in Monitoring HOWTO located at: http://www-1.ibm.com/servers/eserver/clusters/library/csmadm.html

Predefined conditions not shown: The **lscondition** and **lsresponse** commands when run on the management server do not show pre-defined conditions and responses. Follow problem determination section in Monitoring HOWTO located at: http://www-1.ibm.com/servers/eserver/clusters/library/csmadm.html

Commands or file replication fails: CSM commands fail or CFM file replication fails. Follow problem determination section in Overview HOWTO located at: http://www-1.ibm.com/servers/eserver/clusters/library/csmadm.html and FAQs, Hints, and Tips section in Set-Up HOWTO located at: http://www-1.ibm.com/servers/eserver/clusters/library/csmsetup.html

rpower or rconsole commands fail: CSM **rpower** and **rconsole** commands fail. Follow problem determination in Remote Control HOWTO located at: http://www-1.ibm.com/servers/eserver/clusters/library/csmremot.html

GPFS checks

Performance problems: Refer to GPFS Problem Determination and GPFS Performance Whitepapers.

GPFS file system failure: Refer to GPFS problem determination.

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 *syslog*.

Use the **lsnode -Al** command to determine the hostname for the Falcon card and the service processor name associated with the node of interest. Then use telnet or web browser to connect using the hostname for the Falcon card and select options to configure SNMP.

Setting up SNMP alerts from Myrinet

The Myrinet 2000 network in Linux Cluster 1350 is installed with monitoring cards. One can use graphical monitoring program **mute** to monitor the whole network for bad events, all of which are logged and reported by the monitoring cards. You can use an SNMP client or a web browser to access monitoring card information. You can even have monitoring cards notify you of bad events by email.

The following Myrinet software packages are required:

- GM software. This 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, and the GM API, library, and header files (current version is 1.4; version 1.5 is expected soon.).
- m3-dist package. Provides the source for building the SNMP library for the GM layer.
- mute (GUI) tool to monitor the Myrinet network (the name will likely change in the not too distant future).

Order in which the software should be built:

- GM including the mt tools.
- m3-dist (has dependency on GM)
- mute (has dependency on GM and m3-dist)

The README-Linux and mt/README that ships with the GM software, the README that ships with the m3-dist software, and the README that ships with the mute software provide comprehensive details on how to build the respective parts.

Currently m3-dist and mute compile against GM 1.5. With GM 1.4 the SNMP library does not build (m3-dist) and building mute isn't straight forward either. So we recommend building the above software against GM 1.5. (Note that GM 1.5 is not generally available yet but is expected to be released soon.)

All of the above Myrinet software can be obtained from:
http://www.myri.com/scs/index.html (for GM, select the*LANai9* software).

Resetting RSA cards

The RSA card is connected to a remote power control strip. Connect to the remote power control strip connected to the failing RSA card and issue a power off/power on sequence to the RSA card's port.

Checking service processor logs

Use the **lsnode -Al** command to determine the hostname for the Falcon card and the service processor name associated with the node of interest. Then use **telnet** or a web browser to connect using the hostname for the Falcon card and select the view log menu option.

Problem determination references

Linux

Red Hat Linux: Red Hat 7.3 and 8.0 documentation: http://www.redhat.com/docs/manuals/linux/

SuSE Linux: SuSE documentation is included in the downloaded code and is not available on the web. Refer to the information package that came along with the Linux distribution.

CSM

Monitoring HOWTO: http://www-1.ibm.com/servers/eserver/clusters/library/csmadm.html

Overview HOWTO: http://www-1.ibm.com/servers/eserver/clusters/library/csmadm.html

Remote Control HOWTO: http://www-1.ibm.com/servers/eserver/clusters/library/csmremot.html

Set-Up HOWTO: http://www-1.ibm.com/servers/eserver/clusters/library/csmsetup.html

GPFS

Linux IBM General Parallel File System for Linux: http://www-1.ibm.com/servers/eserver/clusters/library/am4pdmst.html

Node

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 bootup. Press F2 to run diagnostics, then F3 to view logfile.

x335 Hardware Maintenance Manual: ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/48p9908.pdf

x335 Installation Guide (includes Troubleshooting): ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/33p2612.pdf

eServer BladeCenter Hardware Maintenance Manual and Troubleshooting Guide: ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/71p9885.pdf

eServer BladeCenter Installation Guide: ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/ga27-4327-00.pdf

x345 Hardware Maintenance Manual and Troubleshooting Guide: ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/48p9718.pdf

x345 Installation Guide: ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/48p9726.pdf

x360 Hardware Maintenance Manual and Troubleshooting Guide: ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/248p2967.pdf x360 Installation Guide: ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/48p9794.pdf

RSA problem determination

IBM eServer xSeries 220, 232, 345 - IBM Remote supervisor adapter user's guide version 5.0: ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/33p2530.pdf

IBM eServer xSeries 335 - Remote supervisor adapter user's guide version 5.0: ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/33p2529.pdf

IBM eServer xSeries 220, 330, 232, 345 - IBM remote supervisor adapter installation guide version 4.0: ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/32p0196.pdf

Storage

There are several ways to check for errors:

- SNMP alerts sent out. The SNMP manager should be set to gather alerts.
- Error indicator light on the device is lit.
- Hard drive indicator light is amber.
- Logging information in the FAStT error log. The log is accessible via **telnet**.
- FAStT can be configured to send messages to the *syslog*.
- Simplicity Storage Manager is a GUID Linux tool that works via Ethernet. This tool can help locate miscablings, throughput issues, and status.

IBM Fibre Channel Hardware Maintenance Manual version 2.0: ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/19k6130.pdf

Hardware maintenance manual for the IBM FAStT 200 Raid Storage Solutions: ftp://ftp.pc.ibm.com/pub/pccbbs/options/06p8884.pdf

ServeRAID-4x Ultra160 SCSI Controllers - Hardware Maintenance Manual: ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers/19k6408.pdf

FAStT Host Adapter Installation and User's Guide: ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers/25p1663.pdf

Cisco switches

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

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

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

Troubleshooting section in Catalyst 2900 XL and Catalyst 3500 XL Software Configuration Guide: http://www.cisco.com/univercd/cc/td/doc/product/lan/c2900xl/29_35wc/sc/

Installation, configuration, technical reference, and release notes for Catalyst 4000 series switches: http://www.cisco.com/en/US/products/hw/switches/ps663/prod_instructions_guides.html

Myrinet

Troubleshooting section in FAQ: http://www.myri.com/scs/GM_FAQ.html

This FAQ has several questions pertaining to Troubleshooting, and is a good starting point.

APC

Troubleshooting for the MasterSwitch: http://www.apcc.com/support/kbase.cfm

Chapter 11. Management, Cluster and Storage Nodes

Contents

The IBM components used for Management, Cluster, and Storage Nodes are shown in Table 20

Table 20. IBM components used for Management, Cluster, and Storage Nodes

Node type	IBM component used
Management Node	• xSeries 345
Cluster Nodes	• xSeries 335 • BladeCenter
Storage Nodes	• xSeries 345 • xSeries 360

While the x335, BladeCenter, x345, and x360 are all high-reliability units, occassionally a component may fail. Two areas that might cause problems are:

1. Disk drives
2. Systemboard

Disk drive failure on the Management Node

The x345 supports hot swapping of hard disks. To replace a failing hard disk on the Management Node:

1. Remove the failing hard disk by sliding it out through the front panel opening.
2. Slide the replacement disk drive into the open slot. The drive will rebuild automatically on a mirrored system. If the system is not mirrored a complete re-install of the Management Node is required. See Chapter 6, “Software installation”, on page 35 for detailed instructions.

Disk drive failure on a Cluster Node

The x335 supports hot swapping of hard disks, but BladeCenter does not. To replace a failing hard disk on an x335:

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

   ```
   installnode x
   ```

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

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

```
installnode x
```

© Copyright IBM Corp. 2003

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

Disk drive failure on a Storage Node

The x345 and x360 support hot swapping of hard disks. To replace a failing hard disk on the Storage Node:

1. Remove the failing hard disk by sliding it out through the front panel opening.
2. Slide the replacement disk drive into the open slot. The drive will rebuild automatically on a mirrored system. If the system is not mirrored enter the following command at the Management Node:

   ```
   installstorage 1
   ```

   and the Storage Node will rebuild.

Systemboard failures

1. Replace the systemboard.
2. Flash the system BIOS to the level used in the installation. Refer to “Software Version Matrix” on page 35 for a listing of the software and firmware levels used in the Cluster 1350.
3. Flash the Diagnostics to match the BIOS level.Refer to “Software Version Matrix” on page 35 for a listing of the software and firmware levels used in the Cluster 1350
4. Flash the onboard ASM to the current level.Refer to “Software Version Matrix” on page 35 for a listing of the software and firmware levels used in the Cluster 1350
5. Make the following configuration settings:
   - **Devices and I/O Ports:** PORT 3F8, IRQ4
   - **Remote Console Redirection:** Enabled, COM1, 9600, 8, None, 1, VT100, Enabled
   - **Boot sequence:** Diskette Drive, CD ROM, Network, Hard Drive 0, Boot Fail Count: DISABLED
   - Set the remote control password if an RSA card is installed in this node (x335, x345, x360 only).
   - If you are replacing an H2O Blade server with a serial port option, ensure that switch 7 in the switchblock is turned **on**.
   - Update the cluster software with the new MAC address associated with the new systemboard or Blade server you installed. Use the CSM command **ifconfig** to get the MAC address of eth0 for the new component. Use the CSM command **chnode <nodename> InstallAdaptorMacaddr=<new MAC (xx:xx:xx:xx:xx:xx)>,Any System** to update the cluster software.

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

Other Cluster Node Problems

x335 problems

In a x335 with an RSA over C2T connection ensure that the Cluster Node at the beginning of the C2T chain has an RSA card, external dongle, and connection to the onboard RSA processor.

BladeCenter problems

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

Take special care when two processors are installed not to pinch the cable under the metal standoff on the inside of the cover.

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

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

Additional Information

Additional hardware maintenance and problem determination information relating to the x335s and x345s was included with the documentation shipped with the Cluster 1350. If you cannot find the manuals use the following links to access online copies:

- **IBM xSeries 335:**
  - *xSeries 335 Hardware Maintenance Manual and Troubleshooting Guide*: ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/48p9908.pdf
- **IBM BladeCenter:**
  - *eServer BladeCenter Hardware Maintenance Manual and Troubleshooting Guide*: ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/71p9885.pdf
- **IBM xSeries 345:**
  - *xSeries 345 Hardware Maintenance Manual and Troubleshooting Guide*: ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/48p9718.pdf
- **IBM xSeries 360:**
  - *xSeries 360 Hardware Maintenance Manual*: ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/24p2967.pdf

Chapter 12. Power problems

Contents

No power to multiple devices

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

No power to an individual device

1. Verify the PDU plugs are securely seated into the Power Out sockets on the Frame Power Blocks.
2. Verify voltage at the Power Out ports on the PDU using a Multimeter. If no power is present replace the PDU (FRU 9N9671). Otherwise, continue with the next step.
3. Swap out the power cable on the failing unit. If power LEDs do not appear on the failing unit, replace the power supply or complete unit if the power supply cannot be replaced.

© Copyright IBM Corp. 2003

Chapter 13. KVM Switch replacement and configuration

Contents

There are twp possible KVM switch options for the Cluster 1350:

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

Replacement of NetBAY 2x8 console switch

To replace the console switch you must remove the rails from the cabinet.

Uncable the console switch

1. Turn power off to the console switch being replaced and remove the power cord.
2. Note the connection positions of signal cables and then disconnect them.

Remove the rails from the cabinet

1. Remove the mounting screws from the front vertical cabinet rails and remove the cover brackets.
2. Remove the mounting screws from the rear vertical cabinet rails.
3. Slide the rails out the rear of the cabinet.

Remove the console switch from the rails

Unscrew the console switch from the rails. There are two holes on each side of the device.

Install the console switch and rails into the cabinet

To install a replacement console switch proceed as follows:

1. Select the location of an added console switch according to the following rules:
   - If a second 1-GB Ethernet Switch is present, or if console switch cable length restrictions require it, install the second console switch in slot 20 of an Expansion cabinet. Nodes to be connected to a console switch cannot have more than one cabinet intervening between their own cabinet and the one where that console switch resides.
   - If a Myrinet switch is present and console switch cable length restrictions allow it, install the second console switch in slot 21 of the Primary cabinet.
   - If no Myrinet switch is present and console switch cable length restrictions allow it, install the second console switch in slot 16 of the Primary cabinet.
2. Mount the switch rails to the console switch.
   a. Mount the rails so that the Cable Management Tab is on the left rail, and the holes in the console switch align with the two smallest holes on each rail, as shown in Figure 10 on page 82.
   b. Use two screws on each side to secure the device.

© Copyright IBM Corp. 2003

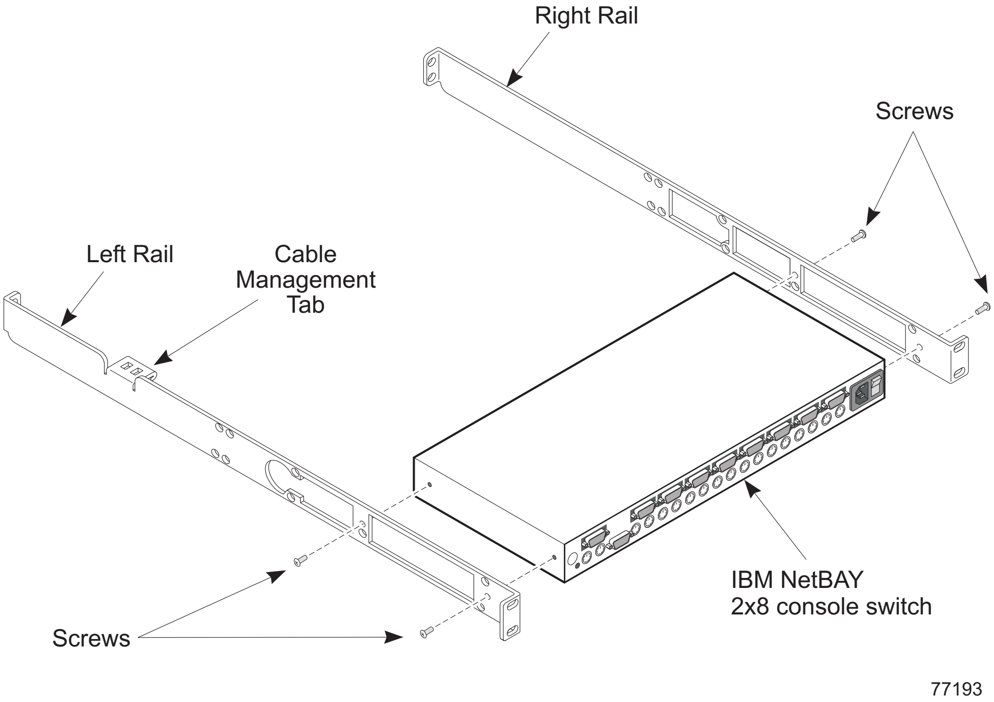

Figure 10. Mount switch rails to the console switch

3. Install the switch rails into the appropriate cabinet and rack slot.
   a. Install clip nuts on the vertical cabinet rails. Install four clip nuts in front and four in back, as shown in the figure below.

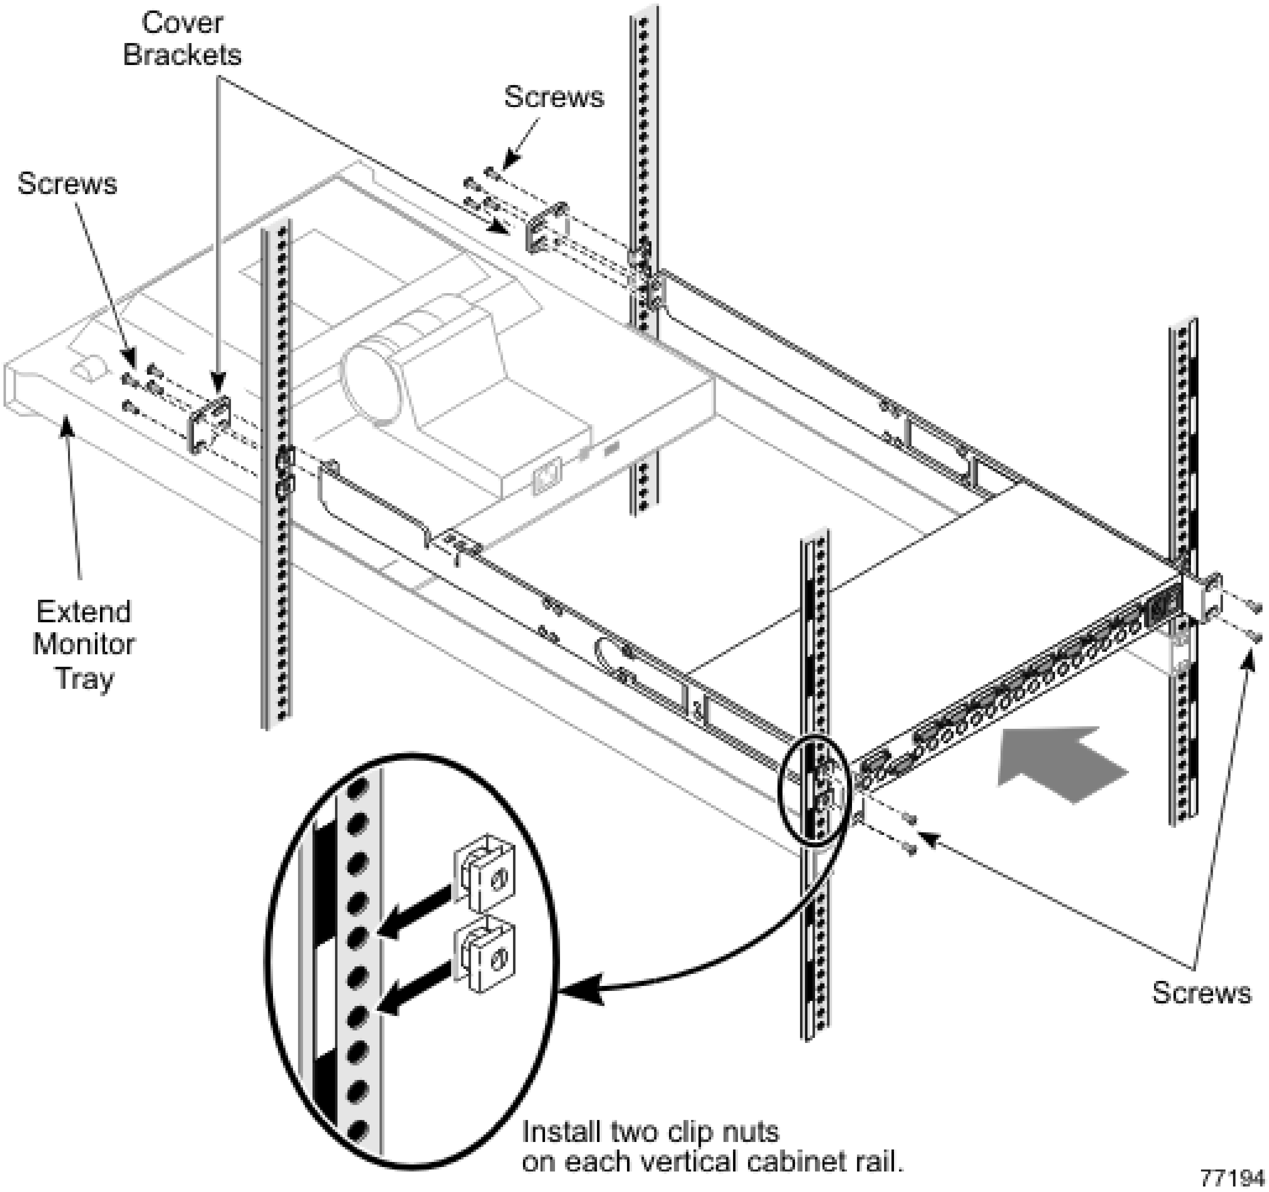

Figure 11. Install switch rails into cabinet from the rear of the cabinet

b. Extend the monitor tray out the front of the cabinet to allow access for installing the rails, as shown in Figure 11.
c. Slide the switch rails into the cabinet from the rear of the cabinet, as shown in Figure 11.
d. Insert the mounting screws into the rear vertical cabinet rails, as shown in Figure 11. Insert two screws on each side. Do not tighten the screws.

   If you have difficulty with neighboring mounting screws of components already installed, loosen these mounting screws then tighten once all of the component rails are installed.
e. At the front of the cabinet, place the cover brackets on the outside of the cabinet vertical rails and insert the mounting screws, as shown in Figure 11. Insert four screws on each side.
f. Tighten the mounting screws in the cabinet vertical rails with a Phillips #3 screwdriver.

4. Cable and power up the console switch.

   a. Connect the console cable.
   b. Connect the power cable and power on the console switch.
   c. Connect the remaining cables in the following order:
      1) Node cables
      2) Video
      3) Mouse
      4) Keyboard
      5) Second console switch (if present)
5. Configure and set up the console switch

Configure and setup the console switch after device replacement

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

Upgrading the console switch FLASH level

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

1. If you haven’t already done so, connect a laptop computer running a terminal emulation program (such as HyperTerminal) to the configuration port on the back panel of the console switch using a RS232 DB9 null modem cable.
2. Set the terminal to 9600 baud, 8 bits, 1 stop bit, no parity and no flow control.
3. Connect the LAN port in the console switch to an Ethernet hub that is also connected to the PC being used as the TFTP server. Launch both the server software and the terminal emulation software.

4. Verify that the console switch is turned on. After approximately 40 seconds, the console switch will send out a message reading : *Avocent AutoView 1000R/2000R Ready_Press any key to continue.* Press any key to access the AutoView 1000R/2000R main menu.
5. Get the IP address of the TFTP server. If using the SolarWinds TFTP server the IP address is shown in the lower right-hand corner of the server's pane. Otherwise you must extract the IP address using the OS tools.
6. Right click on *Network Neighborhood* and select the *Properties* tab.
7. On the *Protocols* tab select *TCP/IP protocol*
8. Select *Properties* and make note of the IP address.
9. If needed, assign the IP address for the console switch:
   a. In the terminal emulation window enter **1** to select *Network Configuration*.
   b. Compare the IP address shown for the console switch to the IP address of the TFTP server. The first three numbers of both IP addresses must be the same, but the last number must be different. If the console switch IP address is incorrect type **3** to select *IP address* and then enter the correct address.
   c. Type **0** to exit the *Network Configuration* menu and follow the prompts on the screen to upgrade the FLASH level on the console switch.

Replacement of NetBAY Advanced Connectivity Technology RCM

Detailed removal, replacement, and configuration information for the RCM is available online at the following URL:
ftp://ftp.pc.ibm.com/pub/pccbbs/pc_servers_pdf/rcm_iug520.pdf

Chapter 14. KVM control

Contents

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

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

Saving the KVM Switch settings

Device settings need to be saved in the KVM Switch's nonvolatile memory (NVRAM) when any of the following occurs:

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

ATTENTION!

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

Perform the following to save the device settings in the KVM Switch's NVRAM:

1. Press **Print Screen**. The OSCAR selection window appears on the display.
2. Press **F2**. The Advanced menu window appears. The highlight is in the Commands menu.
3. Using the **Up** and **Down** arrow keys, move the highlight to Snapshot and press **Enter**. The device settings are now saved to NVRAM.

Connecting components with KVM Switch power on

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

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

Switching between nodes and the console

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

Perform the following to switch between nodes or the console:

1. Press **Print Screen**. The OSCAR selection window appears on the display.

© Copyright IBM Corp. 2003

ATTENTION!
The servers and the console are listed in order by port or by name, depending on the user-definable settings in OSCAR menu attributes.

2. To select a node or the console, do one of the following:
   a. Using the **Up** and **Down** arrow keys to select the node or the console; then press **Enter**.
   b. Press the numeric key that corresponds to the node's port number or the console's port number, then press **Enter**.
   c. Use the mouse to double click on the node or the console you want to select.
3. Press **Escape** to exit OSCAR and to remove the OSCAR selection window from the display. The status flag window returns to the display to indicate the currently connected node or the console.

Security features

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

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

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

Resetting the mouse and keyboard

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

1. Press **Print Screen**. The OSCAR selection window appears on the display.
2. Press **F2**. The Advanced menu window appears. The highlight is in the Commands menu.
3. Using the **Up** and **Down** arrow keys, move the highlight to Reset, and press **Enter**. The mouse and keyboard are now reset.

If the mouse or keyboard are still locked up, you can push the reset button on the back panel to reset the KVM Switch. Pressing the reset button may allow you to recover the device settings without power cycling the node.

Chapter 15. Port server replacement and configuration

Contents

Replacement

The port server can be replaced either online or offline.

To replace the port server, you must remove the tray from the cabinet.

Uncable the port server

1. Obtain the IP address of the port server
2. Power down the port server
3. Note the connection positions of cables to the port server, then disconnect the cables

Remove the tray from the cabinet

To remove the tray from the cabinet, do the following:

1. Remove the mounting screws from the front vertical cabinet rails
2. Retract the sliding bracket at the back of the tray. Retracting the bracket allows you to slide the tray out the rear of the cabinet.
3. Tighten one screw on each side of the sliding bracket to prevent the bracket from extending as you slide the tray from the cabinet
4. Remove the mounting screws from the rear vertical cabinet rails
5. Unplug the power cables and position them out of the way
6. Slide the tray out the rear of the cabinet

Remove the port server from the tray

1. From the bottom of the tray, remove the three screws that secure the port server
2. Slide the port server out of the tray
3. Record the MAC address from the back of the port server

Install the port server and rails

1. Remove all four rubber feet from the base of each port server
2. Mount each port server onto the tray. Mount port server #1 in the right-most position. If the system has a second port server, mount it in the left-most position.
   a. Align the screws on the side of the port server with the alignment holes in the side of the tray, so the screws are visible through the alignment holes.

      The screws on the side of the port server are for positioning only. The securing screws are on the bottom of the port server.
   b. From the bottom of the tray, insert screws into the three holes for each port server. Depending on your assembly environment, it may be easier to flip the port server and tray upside down. Tighten the screws with a Phillips #1 screwdriver.
3. Prepare the tray for installing in the cabinet
   a. Retract the sliding bracket at the back of the tray. Retracting the bracket allows you to slide the tray in from the rear of the cabinet.
   b. Tighten one screw on each side of the sliding bracket to prevent the bracket from extending as you slide the tray in the cabinet

© Copyright IBM Corp. 2003

4. Install the tray into the cabinet into Slot 19
   a. Install clip nuts on the vertical cabinet rails in slot 19. Install four clip nuts in front and four in back.

      The clip nuts may still be on the vertical cabinet rails after the tray was removed.
   b. Slide the tray into the cabinet from the rear of the cabinet.

      If the tray catches as you are sliding it into the cabinet, push up on the tray from underneath.
   c. Insert the mounting screws in the rear vertical cabinet rails. . Do not tighten the screws.

      If you have difficulty with neighboring mounting screws of components already installed, loosen these mounting screws then tighten once all of the component rails are installed.
   d. At the front of the cabinet, extend the sliding brackets so the holes line up with the front vertical cabinet rails. Insert the mounting screws in the front vertical cabinet rails.
   e. Tighten the mounting screws in the cabinet vertical rails with a Phillips #3 screwdriver
   f. Tighten the sliding bracket screws in the tray
5. Attach cables and then power up the port server.
6. Configure and set up the port server

Configuration and setup after device replacement

If you can successfully **ping** the InREACH port server no further action is needed, the port server has been properly configured. If you cannot **ping** the port server, then take the following steps to configure the device:

1. Connect a laptop computer running a terminal emulation program (such as HyperTerminal) to the command port on the back panel of the port server using a DB9 to RJ45 Serial cable.
2. Set the terminal to 9600 baud, 8 bits, 1 stop bit, no parity and no flow control.
3. Insert the PCMCIA flash card that came with the Cluster 1350 into the slot in the front of the port server.
4. Attach a serial terminal to the command port. The default command port is the last port, either port 20 or port 40 depending on the size of the port server.
5. Power on the port server and press **Enter** until the ***Login>*** prompt appears.
6. At the ***Login>*** prompt type *access*. No characters will appear. Press *Enter*.
7. At the ***Username>*** prompt type *system*
8. At the ***IN-Reach>*** prompt type *set priv*
9. At the ***Password>*** prompt type *system*
10. At the ***IN-Reach_Priv>*** prompt type *show ip* to see the current network settings.
11. Type *define ip address xxx.xxx.xxx.xxx* to set the IP address
12. Type *define ip primary gateway address xxx.xxx.xxx.xxx* to set the gateway address
13. Type *define ip subnet mask xxx.xxx.xxx.xxx* to set the subnet mask
14. Type *init delay 0* to save the configuration and restart the port server.

If the port server did not already have an IP address, it may need further configuration so the serial ports can operate properly. Take the following steps:

1. Telnet to the IP address assigned to the port server.
2. At the ***Login>*** prompt type *access*.
3. At the ***Username>*** prompt type *system*
4. At the ***IN-Reach>*** prompt type *set priv*

5. At the ***Password>*** prompt type *system*
6. Define the ports by entering the following at the ***IN-Reach_Priv>*** prompt. Do not try to define ports 21–40 on a port server with only 20 ports.

   ```
   IN-Reach_Priv>define port 1-20 access remote
   IN-Reach_Priv>define port 21-40 access remote
   IN-Reach_Priv>define port 1-20 flow control enable
   IN-Reach_Priv>define port 21-40 flow control enable
   IN-Reach_Priv>define port 1-20 speed 9600
   IN-Reach_Priv>define port 21-40 speed 9600
   IN-Reach_Priv>define port 1-20 que disable
   IN-Reach_Priv>define port 21-40 que disable
   IN-Reach_Priv>lo port 1-20
   IN-Reach_Priv>lo port 21-40
   IN-Reach_Priv>init delay 0
   ```

   The last command will cause the port server to save any configuration changes and restart. The port server should now be fully operational.

Chapter 16. Cisco 10/100 Switch replacement and configuration

Contents

Replacement of the 24-port switch

To replace the Cisco 24-port 10/100 Switch, you must remove the rails from the cabinet.

Remove the rails from the cabinet

1. Remove the mounting screws from the front vertical cabinet rails and remove the cover bracket.
2. Remove the mounting screws from the rear vertical cabinet rails.
3. Slide the rails out the rear of the cabinet.

Remove the Cisco 24-port 10/100 Switch from the rails

Unscrew the Cisco 10/100 Switch from the rails. There are four holes on each side of the devices.

Install the Cisco 24-port 10/100 Switch and rails

1. Mount the switch rails to the Cisco 24-Port 10/100 Switch. Mount the rails so the Cable Management Tab is on the left rail, as shown in the figure below. Use four screws on each side. Use the set of two large holes at the front of the rail and the set of two large holes that are the third set from the back of the rail.

© Copyright IBM Corp. 2003

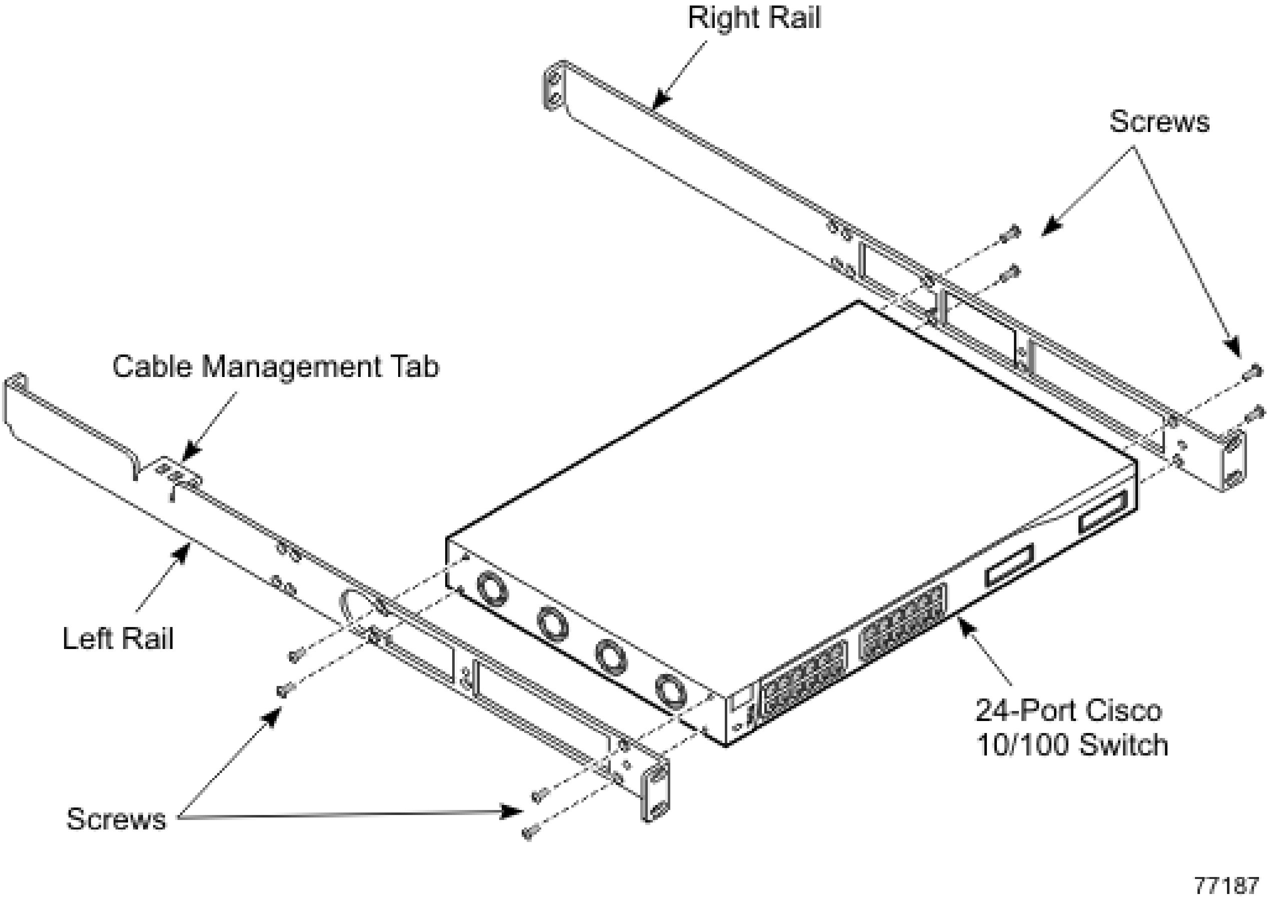

Figure 12. Mount switch rails to the Cisco 24-port 10/100 Switch

2. Install the switch rails into the cabinet in Slot 17.
   a. Install clip nuts on the vertical cabinet rails in Slot 17. Install four clip nuts in front and four in back, as shown in the figure below.

      The clip nuts may still be on the vertical cabinet rails after the rails were removed.

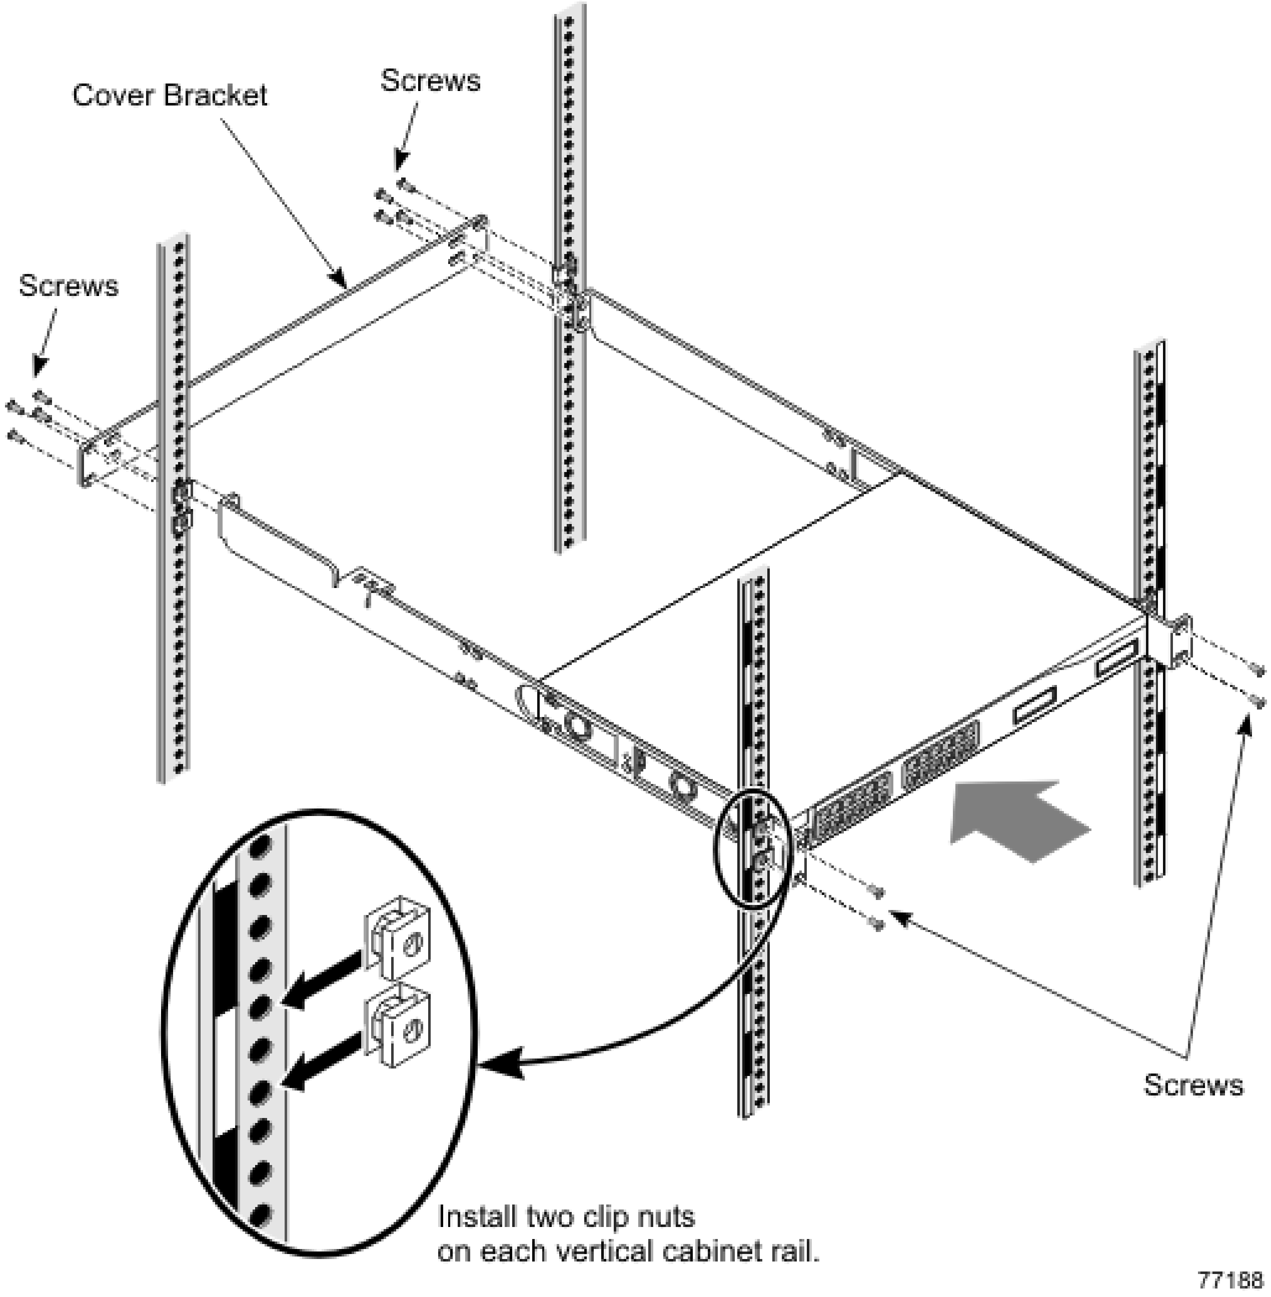

Figure 13. Install switch rails from the rear of the cabinet

b. Slide the switch rails into the cabinet from the rear of the cabinet, as shown in Figure 13. Slide the rails part way in and route the power cables and Ethernet cables from the front of the cabinet. Route the cable through the Cable Management Tab shown in Figure 12 on page 94.

c. Insert the mounting screws in the rear vertical cabinet rails, as shown in Figure 13. Insert two on each side. Do not tighten the screws.

If you have difficulty with neighboring mounting screws of components already installed, loosen these mounting screws then tighten once all of the component rails are installed.

   d. At the front of the cabinet, place the long cover bracket on the outside of the cabinet vertical rails and insert the mounting screws, as shown in Figure 13 on page 95. Insert four screws on each side.
   e. Tighten the mounting screws in the cabinet vertical rails with a Phillips #3 screwdriver.
3. Attach cables and then power up the Cisco 10/100 Switch.
4. Configure and setup the Cisco 10/100 Switch. Refer to “Configuration and setup after device replacement” on page 99.

Replacement of the 48-port switch

To replace the Cisco 48-port 10/100 Switch, you must remove the rails from the cabinet.

Remove the rails from the cabinet

To remove the tray from the cabinet, do the following:

1. Remove the mounting screws from the front vertical cabinet rails and remove the cover bracket.
2. Remove the mounting screws from the rear vertical cabinet rails.
3. Slide the rails out the rear of the cabinet.

Remove the Cisco 48-Port 10/100 Switch from the rails

Unscrew the Cisco 10/100 Switch from the rails. There are four holes on each side of the devices.

Install the Cisco 48-Port 10/100 Switch and rails

1. Mount the switch rails to the 48-Port Cisco 10/100 Switch. Mount the rails so the Cable Management Tab is on the left rail, as shown in the figure below. Use four screws on each side. Use the set of two large holes at the front of the rail and the set of two large holes that are the second set from the back of the rail.

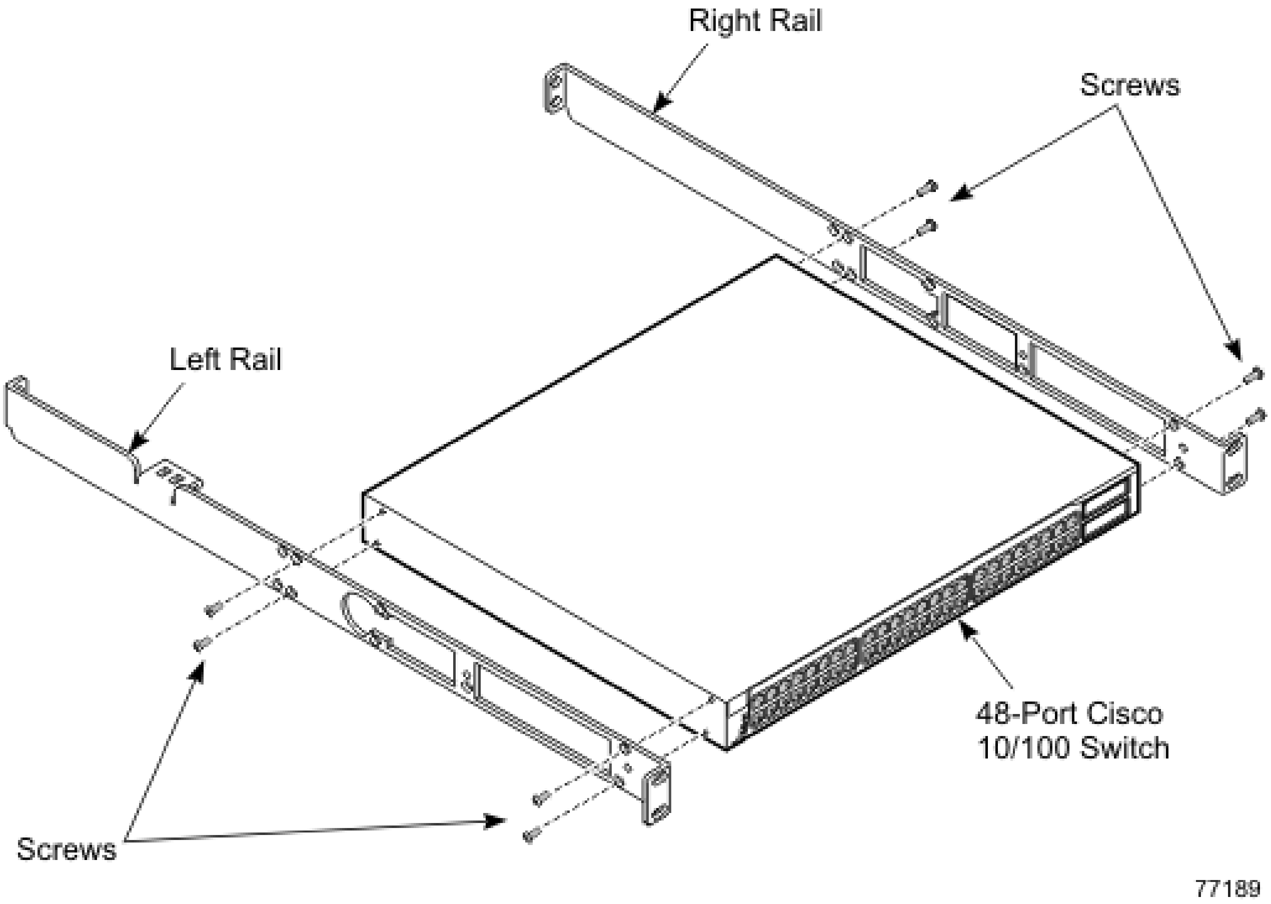

Figure 14. Mount switch rails to the Cisco 48-port 10/100 Switch

2. Install the switch rails into the cabinet.
   a. Install clip nuts on the vertical cabinet rails. Install four clip nuts in front and four in back, as shown in the figure below.

      The clip nuts may still be on the vertical cabinet rails after the rails were removed.

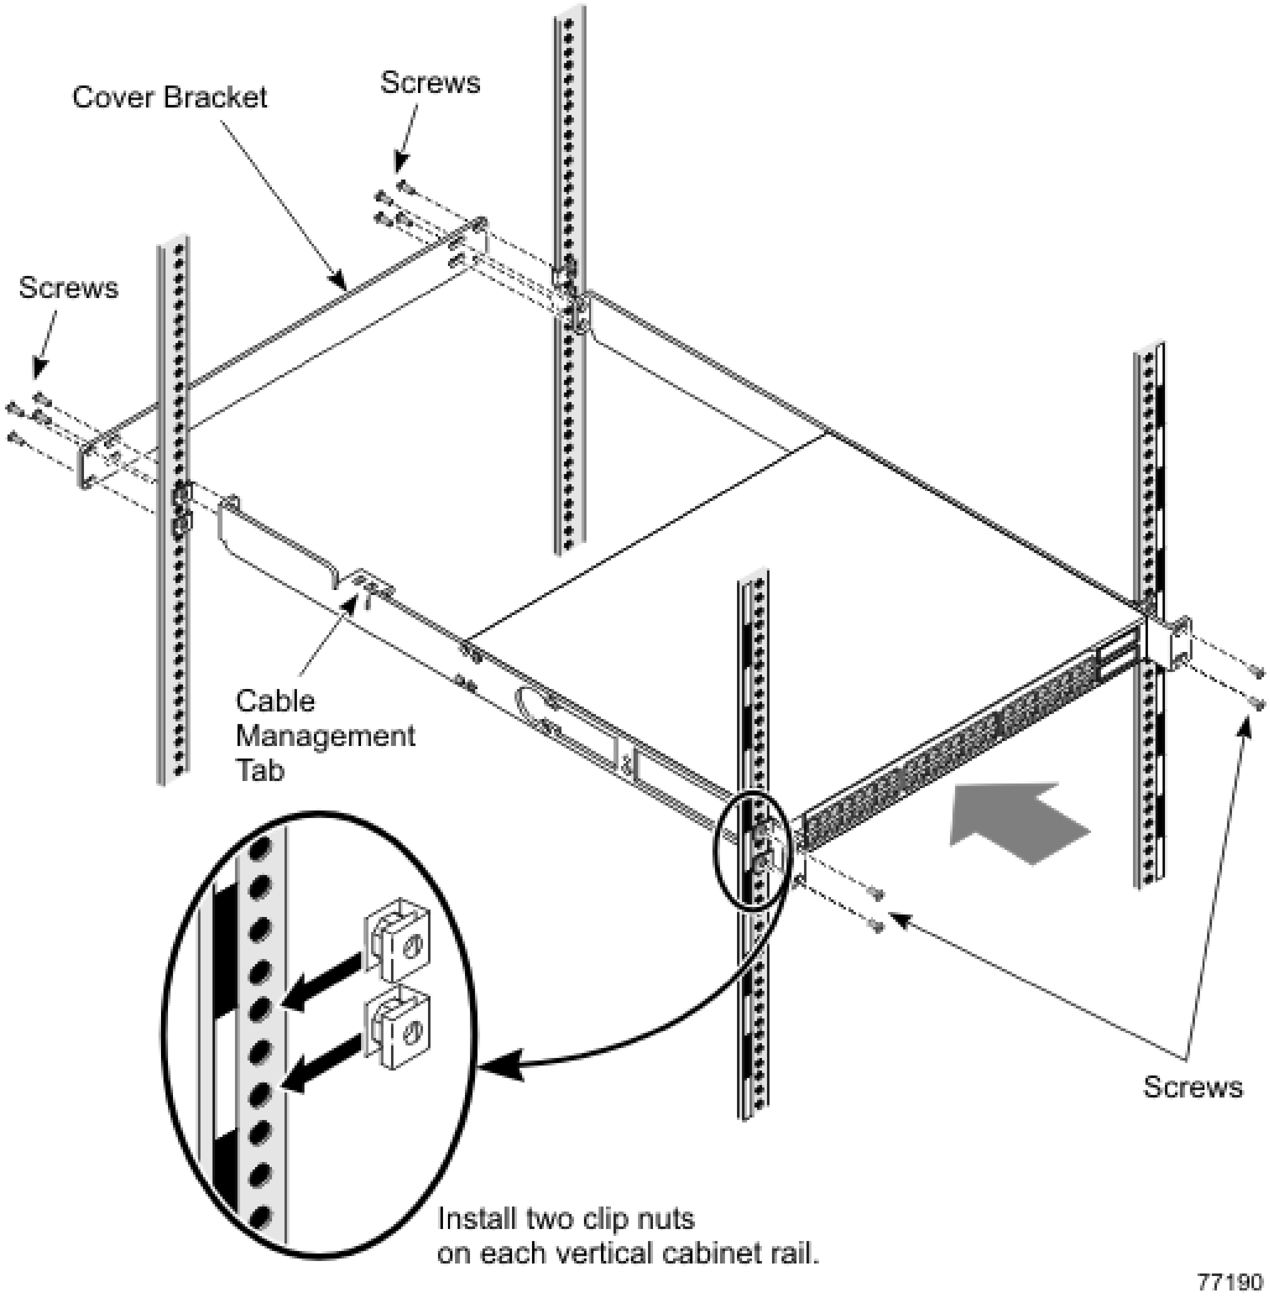

Figure 15. Install switch rails into cabinet from the rear of the cabinet

b. Slide the switch rails into the cabinet from the rear of the cabinet, as shown in Figure 15. Slide the rails part way in and route the power cables and Ethernet cables from the front of the cabinet. Route the cable through the Cable Management Tab shown in Figure 14 on page 97.

c. Insert the mounting screws in the rear vertical cabinet rails, as shown in Figure 15. Insert two screws on each side. Do not tighten the screws.

   If you have difficulty with neighboring mounting screws of components already installed, loosen these mounting screws then tighten once all of the component rails are installed.

   d. At the front of the cabinet, place the long cover bracket on the outside of the cabinet vertical rails and insert the mounting screws, as shown in Figure 15 on page 98. Insert four screws on each side.
   e. Tighten the mounting screws in the cabinet vertical rails with a Phillips #3 screwdriver.
3. Attach cables and then power up the Cisco 10/100 Switch.
4. Configure and setup the Cisco 10/100 Switch. Refer to "Configuration and setup after device replacement".

Configuration and setup after device replacement

To set up the new 10/100 Switch you will need the following:

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

Take the following steps:

1. Connect the RJ45 side of the serial cable to the port on the front of the Cisco switch marked ″CONSOLE.″
2. Connect the other end of the cable to the laptop computer.
3. Start the Hyper Terminal application. Configure the terminal to **9600, 8, N, 1, No Flow Control and VT100 Emulation**.
4. At the command prompt in the terminal emulation window enter **enable**. This will put you in administrative mode.
5. At the prompt type **ibm** and hit enter. The prompt will change from a **>** to a **#** to indicate you are in administrative mode.
6. Enter**show run** to show the current configuration information. Make note of the current settings.

   You need the following information to set up the new switch:
   - Switch IP address
   - IP mask
   - Default gateway IP address
   - Switch host name
   - Cluster name
7. At the # prompt enter **configure terminal**.
8. Enter **interface vlan1**.
9. Enter **ip address 172.xxx.xxx.xxx 255.255.xxx.xxx**.
10. Enter **exit**.
11. Enter **ip default-gagteway 172.xxx.xxx.xxx**.
12. Enter **end**.
13. Enter **show run-config** to verify the IP settings.
14. Enter **copy running-config startup-config**.
15. Logoff from the session.

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

After completing the First Time Setup section in the Quick Start Guide, save it to the startup file so the switch can be rebooted without losing the setup. At the telnet prompt, enter the command:

copy run start

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

ATTENTION!

There is an SNMP vulnerability for various versions of switch firmware. Refer to http://www.cisco.com/warp/public/707/cisco-malformed-snmp-msgs-pub.shtml for specific firmware patches to download.

Setup troubleshooting

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

- On the PC, use the command: **ipconfig**
- On the switch, use the command: **show running**

Additional information

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

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

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

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

Chapter 17. Cisco Gigabit Switch replacement and configuration

Contents

Installation procedure

To replace the Cisco Gigabit Switch, you must remove the rails from the cabinet.

Remove the rails from the cabinet

1. Remove the mounting screws from the front vertical cabinet rails and remove the cover bracket.
2. Remove the mounting screws from the rear vertical cabinet rails.
3. Slide the rails out the rear of the cabinet.

Remove the Cisco Gigabit Switch from the rails

Unscrew the Cisco Gigabit Switch from the rails. There are four holes on each side of the devices.

Install the Cisco Gigabit Switch and rails into the cabinet

1. Mount the switch rails to the Cisco Gigabit Switch.
   a. Mount the rails so the Cable Management Tab is on the left rail, and the holes in the Gigabit Switch align with the set of two large holes at the front of the rail and the set of two large holes closest to the back of the rail, as shown in the figure below.

© Copyright IBM Corp. 2003

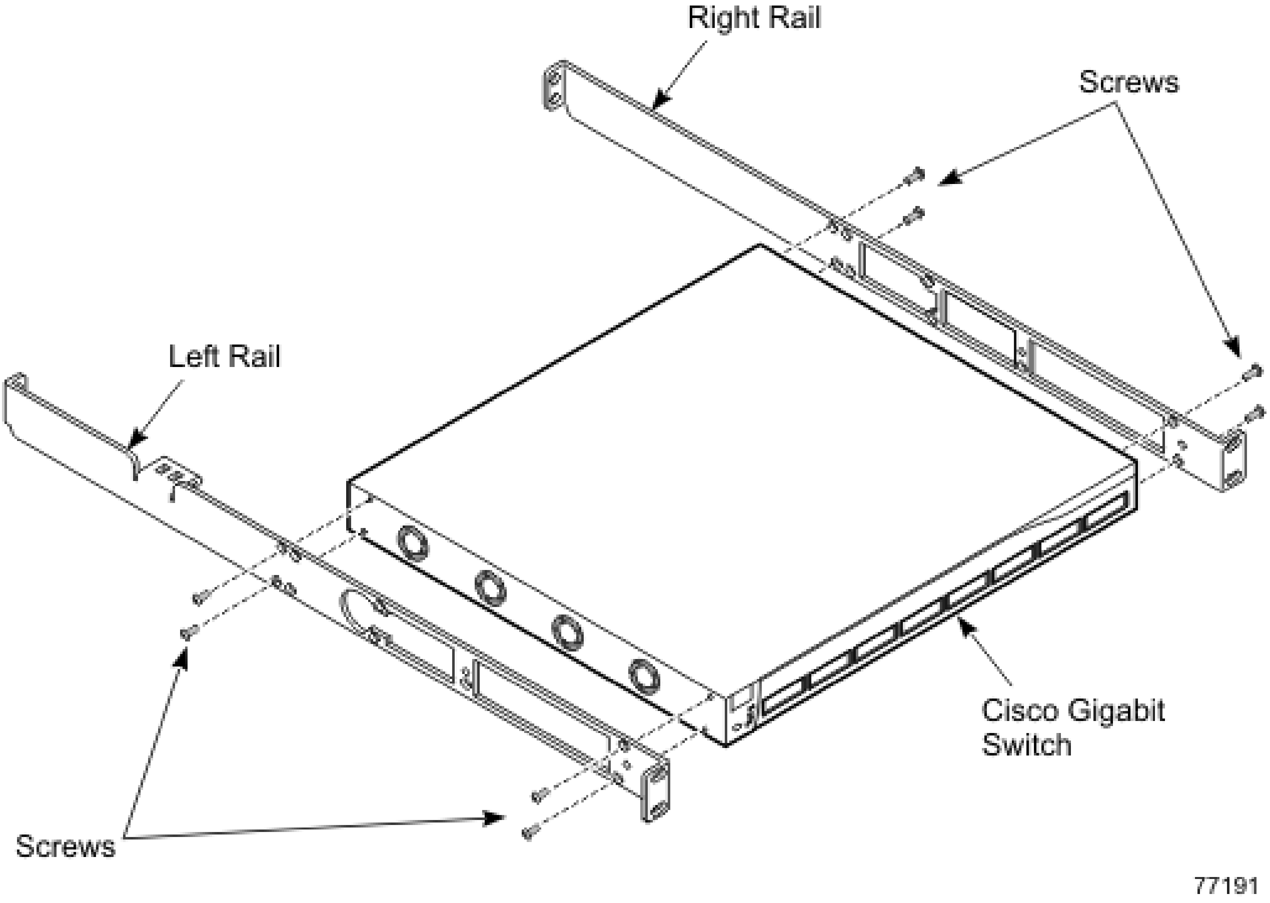

Figure 16. Mount switch rails to the Cisco Gigabit Switch

   b. Use four screws on each side to secure the device.

2. Install the switch rails into the cabinet in Slot 20.

   a. Install clip nuts on the vertical cabinet rails in Slot 20. Install four clip nuts in front and four in back, as shown in the figure below.

      The clip nuts may still be on the vertical cabinet rails after the rails were removed.

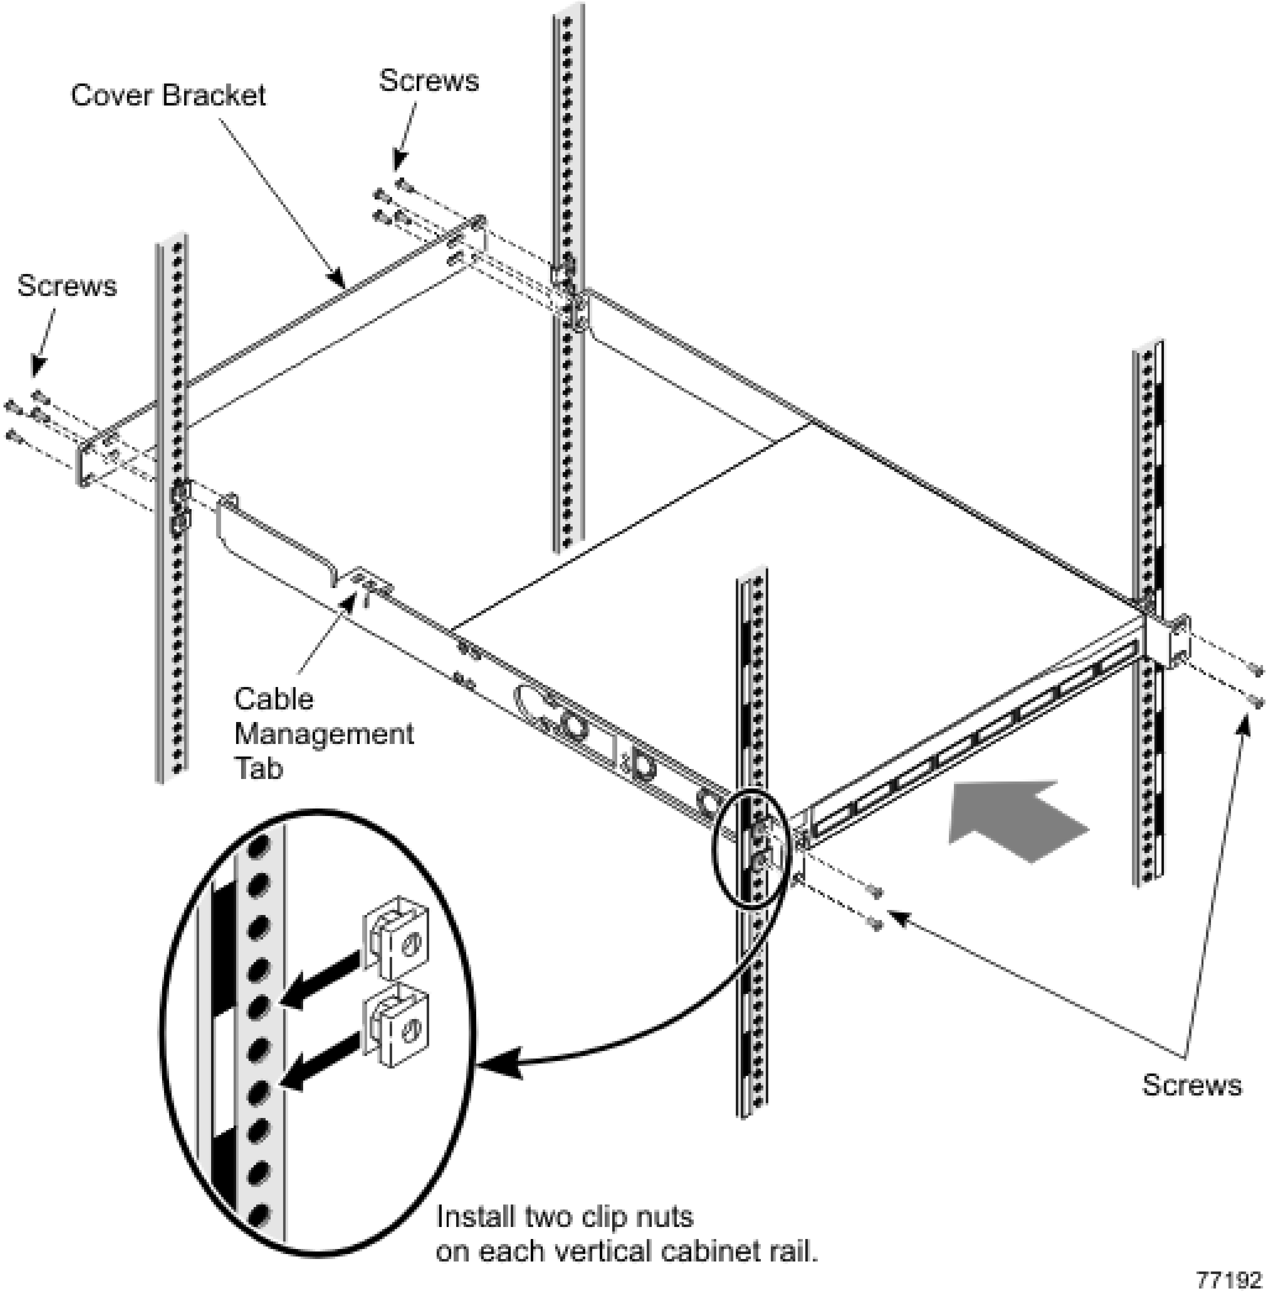

Figure 17. Install switch rails into cabinet from the rear of the cabinet

b. Slide the switch rails into the cabinet from the rear of the cabinet, as shown in Figure 17. Slide the tray part way in and route the power cables and Ethernet cables from the front of the cabinet. Route the cable through the Cable Management Tab shown in Figure 16 on page 102.

c. Insert the mounting screws in the rear vertical cabinet rails, as shown in Figure 17. Insert two on each side. Do not tighten the screws.

   If you have difficulty with neighboring mounting screws of components already installed, loosen these mounting screws then tighten once all of the component rails are installed.

d. At the front of the cabinet, place the long cover bracket on the outside of the cabinet vertical rails and insert the mounting screws, as shown in Figure 17 on page 103. Insert four screws on each side.
e. Tighten the mounting screws in the cabinet vertical rails with a Phillips #3 screwdriver.

3. Attach cables and power up the Gigabit Switch.
4. Perform the "Configure and setup after device replacement" procedure for the Gigabit Switch.

Configure and setup after device replacement

To set up the new Gigabit Switch you will need the following:

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

Take the following steps:

1. Connect the RJ45 side of the serial cable to the port on the front of the Cisco switch marked ″CONSOLE.″
2. Connect the other end of the cable to the laptop computer.
3. Start the Hyper Terminal application. Configure the terminal to **9600, 8, N, 1, No Flow Control and VT100 Emulation**.
4. At the command prompt in the terminal emulation window enter **enable**. This will put you in administrative mode.
5. At the prompt enter **ibm** and hit enter. The prompt will change from a **>** to a **#** to indicate you are in administrative mode.
6. Enter **show run** to show the current configuration information. Make note of the current settings and then logoff from the session.

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

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

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

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

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

ATTENTION!

There is an SNMP vulnerability for various versions of switch firmware. Refer to: http://www.cisco.com/warp/public/707/cisco-malformed-snmp-msgs-pub.shtml for specific firmware patches to download.

Setup troubleshooting

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

- On the PC use the command: **ipconfig**
- On the switch use the command: **show running**

Nodes on the same VLAN can communicate via **ping/telnet**. They cannot communicate to nodes on different VLANs. To verify VLANs:

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

Additional information

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

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

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

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

Chapter 18. Cisco 4000 Series switch replacement

Contents

Installation, removal, replacement, and troubleshooting procedures

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

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

Additionally, IBM has included with each Cisco 4000 Series switch a specially designed thermal duct to ensure proper airflow around the switch. Figure 18 shows an exploded view of the thermal duct and how it fits within the cabinet and attaches to the switch.

Figure 18. Thermal duct used with Cisco 4000 switches

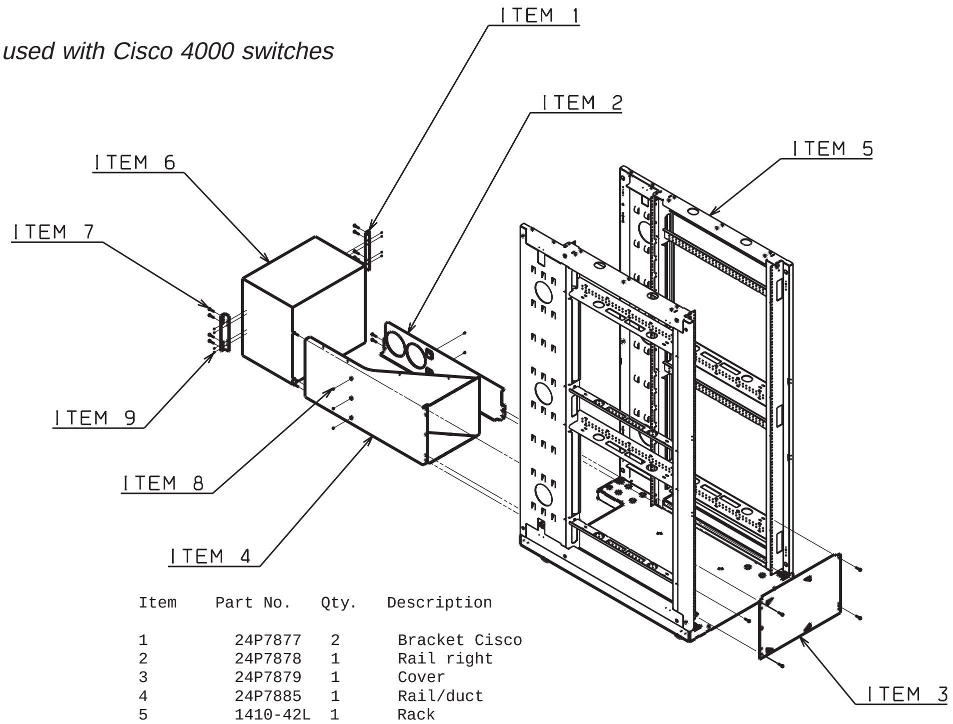

Item	Part No.	Qty.	Description
1	24P7877	2	Bracket Cisco
2	24P7878	1	Rail right
3	24P7879	1	Cover
4	24P7885	1	Rail/duct
5	1410-42L	1	Rack
6	Ref. only	1	Cisco 400x switch
7	12J5289	1	M6 Hex flange screw
8	1621811	5	M4 Hex flange screw
9	N/A	6	M4 Flat head screw

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

Additional information

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

Chapter 19. Myrinet 2000

Contents

The 2-Gbit Myrinet Switch is a customer option that provides high-speed communication between the Storage Nodes and Cluster Nodes, over an optical cable. It requires a Myrinet Switch chassis in the Primary Cabinet and a Myrinet PCI adapter in each Storage Node and Cluster Node.

Myrinet PCI board

The Myrinet PCI board resides in the Cluster (x335) and Storage (x345, x360) Nodes. Use the installation procedures in the appropriate server's documentation to replace a Myrinet PCI board.

The GM software for running the Myrinet board should already reside on the Cluster and Storage Nodes, so no new installation of software should be required when a Myrinet PCI board is replaced.

Myrinet switch chassis

The Myrinet Switch contains the following replaceable components:

8-port line card
Provides the connections to the Storage and Cluster Nodes. The line cards plug into slots in the Switch chassis.

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

Blower module
Cools the Myrinet Switch Chassis

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

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

Table 21. Myrinet Chassis Sizes and Capacities

Slots in Switch	Line Cards	Nodes Supported	EIA Slots Consumed
5	1-4	4-32	4
9	1-8	4-64	6
17	1-16	4-128	10

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

1. Ensure that the cluster is not running critical applications.
2. If the optical cables connected to the Switch are not labeled, place labels on the cables so they can be located to their respective connectors when the new chassis is installed.

© Copyright IBM Corp. 2003

3. Disconnect the optical cables from the connectors on the Myrinet Switch. You do not need to power down or deconfigure the Switch before doing this. **Be sure to install dust caps on all the connectors after the cables are removed.**
4. Disconnect the power cord from the Myrinet Switch. This powers down the Switch.
5. Remove the rack-mount screws from the chassis; then remove the chassis from the rack.
6. Install the new chassis and fasten the rack-mount screws.
7. Connect the optical cables to the connectors on the Switch. **Save the dust caps for future use.**
8. Connect the power cord to the Myrinet Switch. This powers up the Switch.

Configure and setup after device replacement

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

IBM Customer Support personnel will update the firmware if necessary.

Additional information

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

Chapter 20. Power Management Module replacement and configuration

Contents

Replacement

The current Power Management Module is the APC MasterSwitch Power Distribution Unit, Model AP9212. We call it the Power Management Module in order to avoid confusion with the IBM Netfinity Power Distribution Unit, the fourteen-outlet power distribution bars that fit into sidepockets on the mounting rack and plug into the main power supply for the site. In this document, 'PDU' will refer exclusively to the power distribution bars.

To replace the Power Management Module, you must remove the tray from the cabinet.

Uncable the Power Management Module

1. Note the positions of all cable connections to the Power Management Module
2. Obtain the IP address of the Power Management Module
3. Unplug the Power Management Module power cable from the PDU
4. Unplug the remote power cables
5. Unplug the cables to the port server

Remove the tray from the cabinet

1. Remove the mounting screws from the front vertical cabinet rails
2. Retract the sliding bracket at the back of the tray. Retracting the bracket allows you to slide the tray out the rear of the cabinet. The figure below shows the location of the sliding bracket, but shows them being extended.

© Copyright IBM Corp. 2003

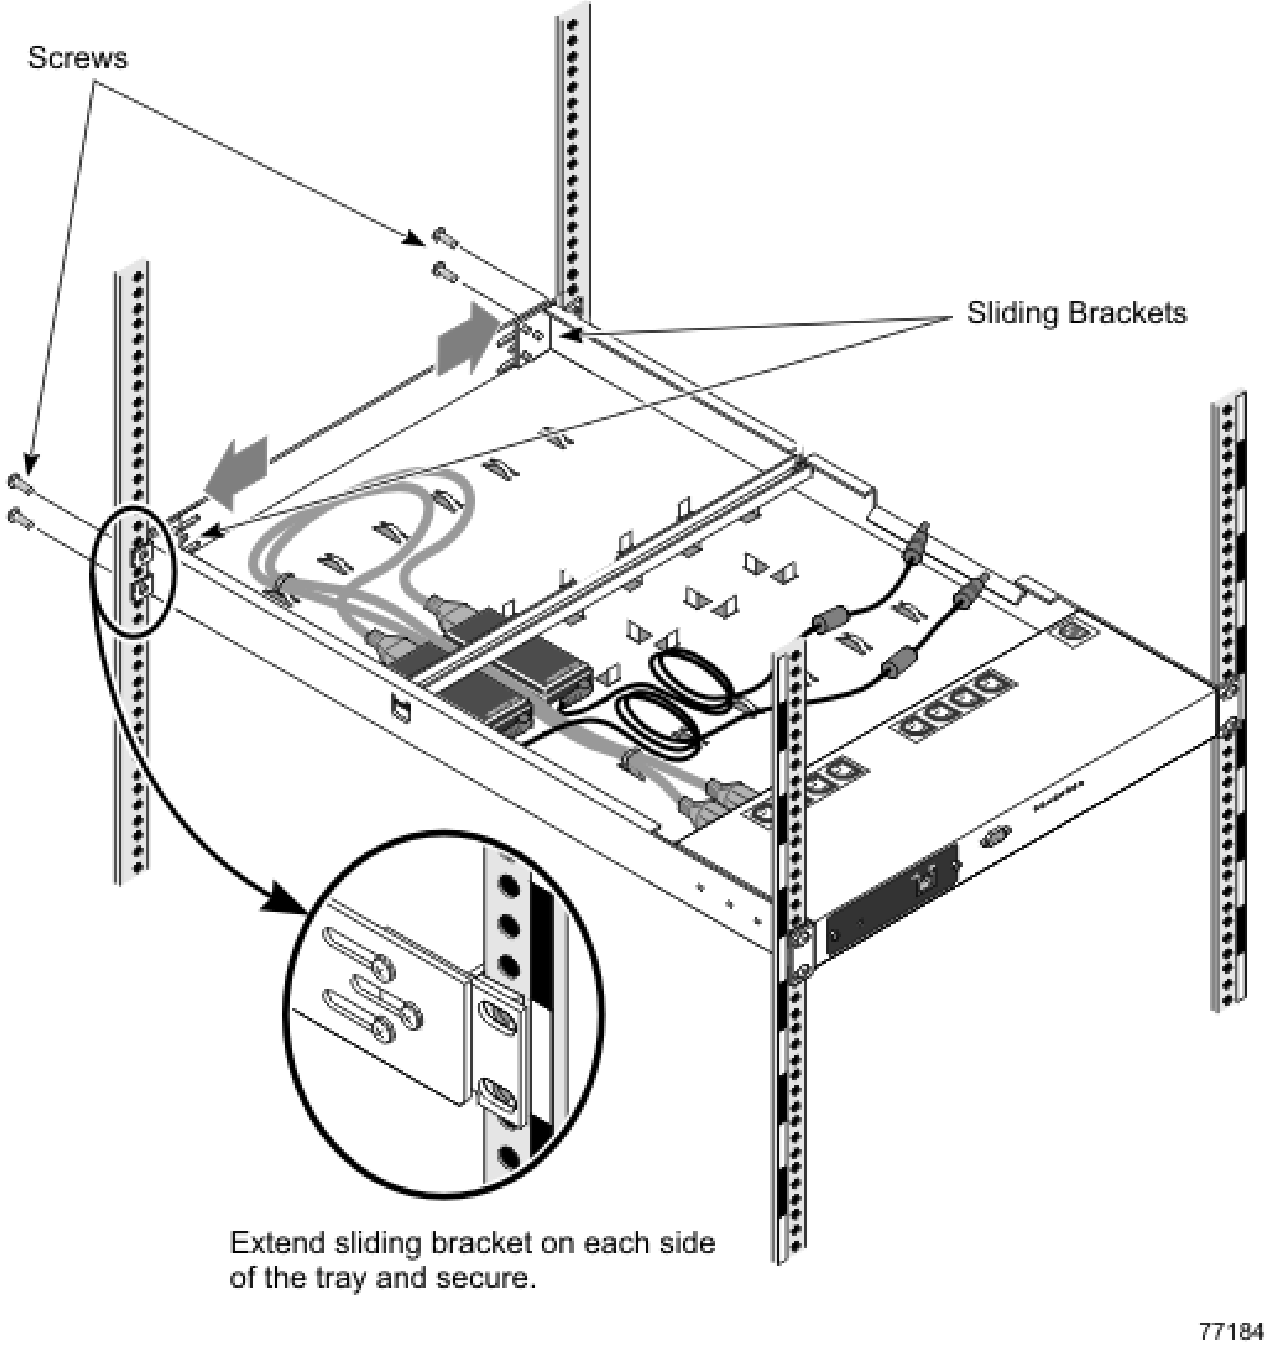

Figure 19. Retract sliding bracket at the back of the tray (front of the cabinet)

3. Tighten one screw on each side of the sliding bracket to prevent the bracket from extending as you slide the tray from the cabinet
4. Remove the mounting screws from the rear vertical cabinet rails
5. Slide the tray out the rear of the cabinet

Remove the Power Management Module from the tray

1. Unplug the power cables from the power ″bricks″ that lie on the tray behind the Power Management Module

2. Unscrew the Power Management Module from the tray. The two mounting screws are located in the outermost of the four holes on each side of the device.
3. Slide the Power Management Module out of the tray

Install the Power Management Module and tray into the cabinet

1. If adding a Power Management Module, determine the cabinet location and rack position of the module.
2. Mount the Power Management Module on the tray
   a. Slide the Power Management Module onto the tray so that the two outside holes on each side of the Power Management Module align with the two outer-most holes on each side of the tray, as shown in Figure 20.

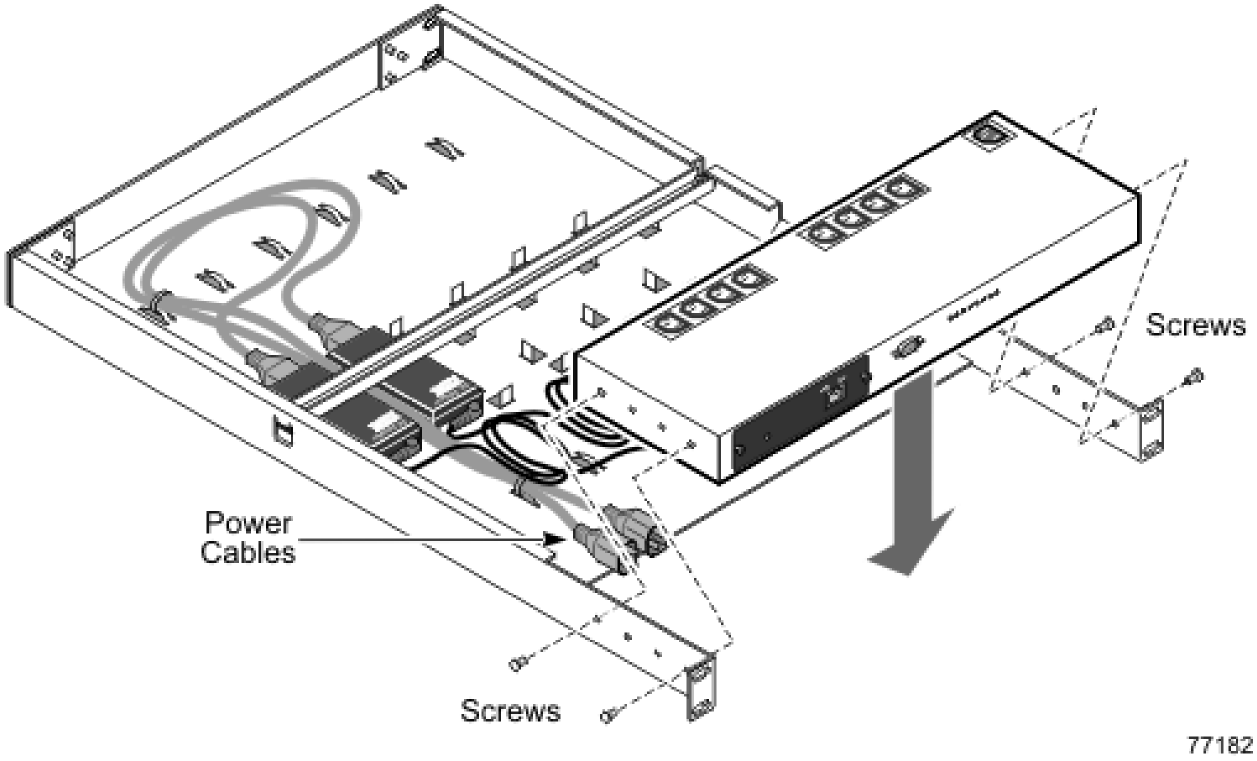

Figure 20. Mount Power Management Module onto Power Management Module tray

   b. Insert two screws on each side of the tray in the two outermost holes. Use the screws that are packaged with the Power Management Module braces. (You will not use the braces, but you need to use the screws.)
   c. Plug the power cables into the Power Management Module. Route the remote power cables for the Nodes out through the cable management slot on the side of the tray, as shown in Figure 21 on page 114.

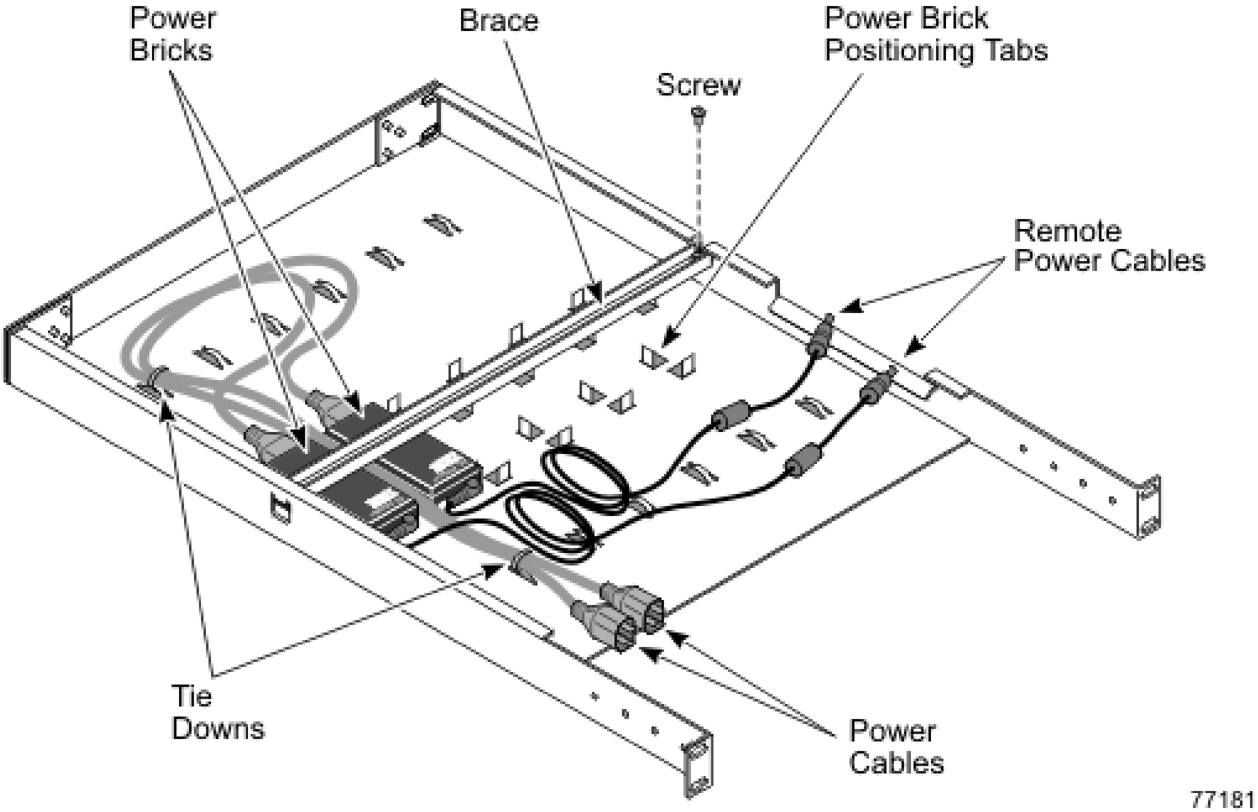

Figure 21. Mount power bricks onto Power Management Module tray

   d. Secure the power cables to the tie-downs in the tray using cable ties, as shown in Figure 21

3. Prepare the Power Management Module tray for installing in the cabinet
   a. Retract the sliding bracket at the back of the tray. Retracting the bracket allows you to slide the tray in from the rear of the cabinet
   b. Tighten one screw on each side of the sliding bracket to prevent the bracket from extending as you slide the tray in the cabinet
   c. Install the clip nuts on the cabinet vertical rails in Slot 18. The clip nuts may still be on the vertical rails after the tray was removed. Install four in front and four in back, as shown in Figure 22 on page 115.

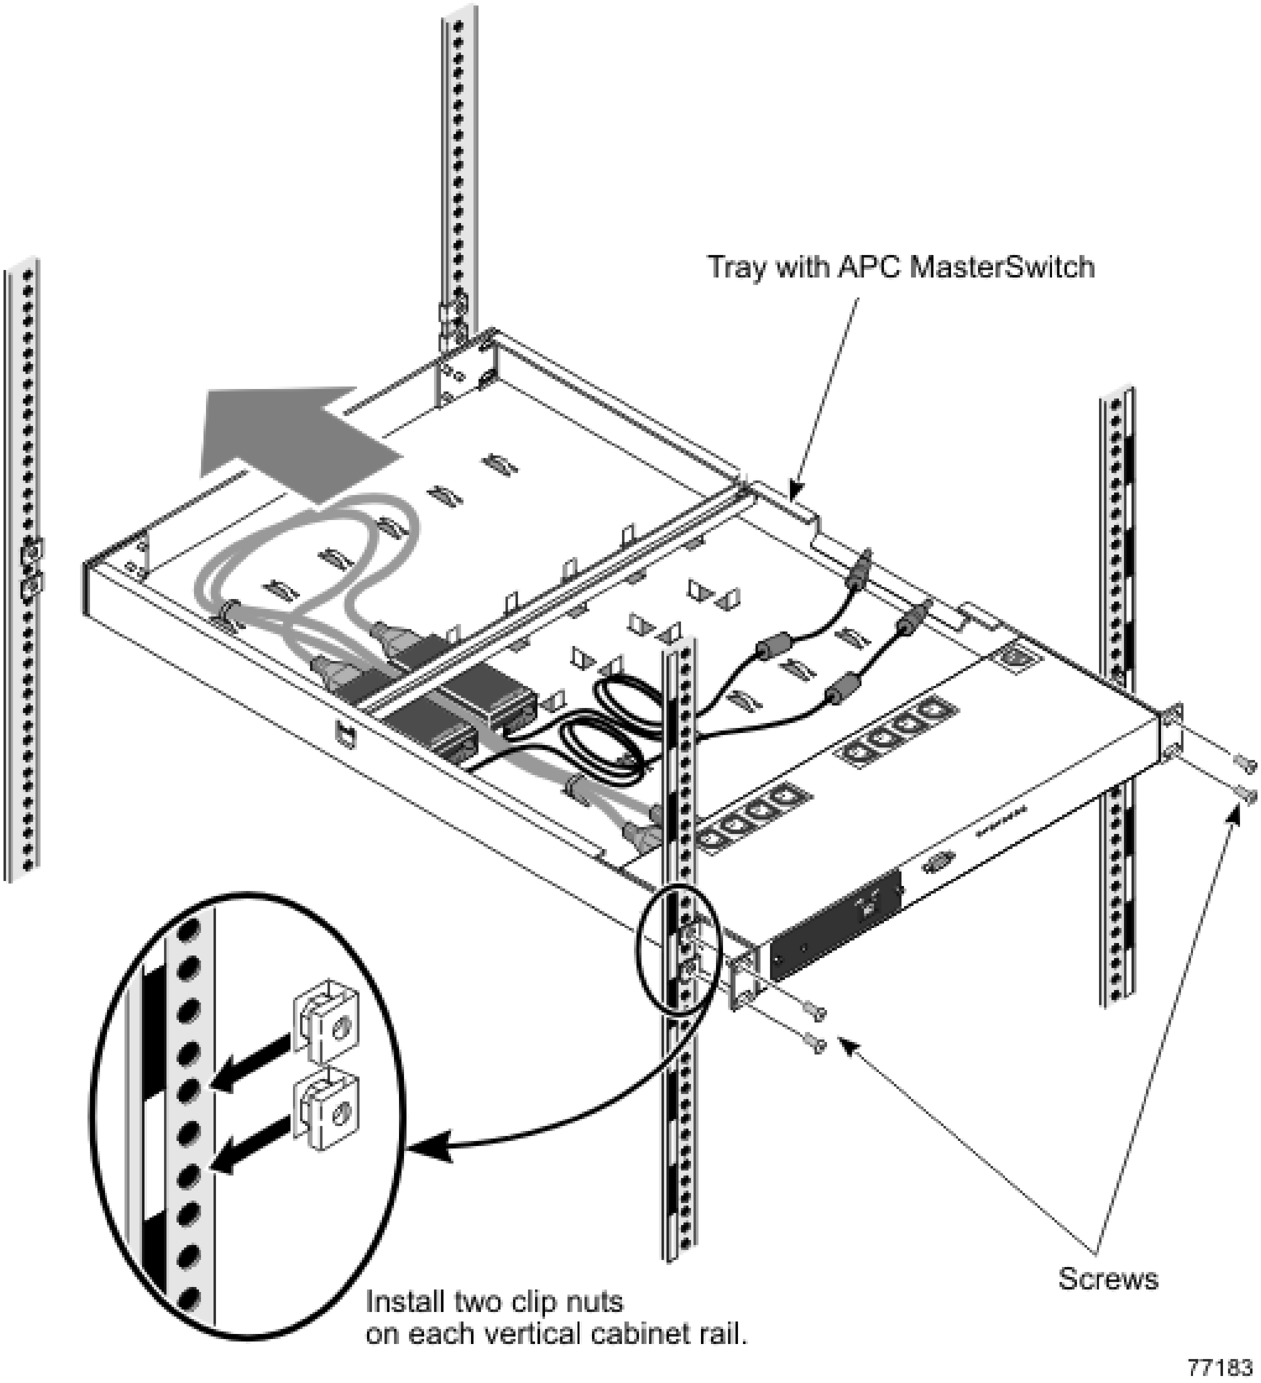

Figure 22. Install Power Management Module tray into cabinet from the rear of the cabinet, part 1

4. Mount the Power Management Module tray into the cabinet in Slot 18.
   a. Slide the tray into the cabinet from the rear of the cabinet, as shown in Figure 22. Slide the tray partially in and route the remote power cables (which go to the RSA cards in the Nodes) and the cable to the PDU out through the cable management slot in the tray.

      If the tray catches as you are sliding it into the cabinet, push up on the tray from underneath.
   b. Insert the mounting screws into the rear vertical cabinet rails, as shown in Figure 22. Do not tighten the screws.

c. If you have difficulty with neighboring mounting screws of components already installed, loosen these mounting screws, then tighten them once all of the component rails are installed

d. At the front of the cabinet, extend the sliding brackets so that the holes line up with the front vertical cabinet rails, as shown in the figure below. Insert the mounting screws in the front vertical cabinet rails.

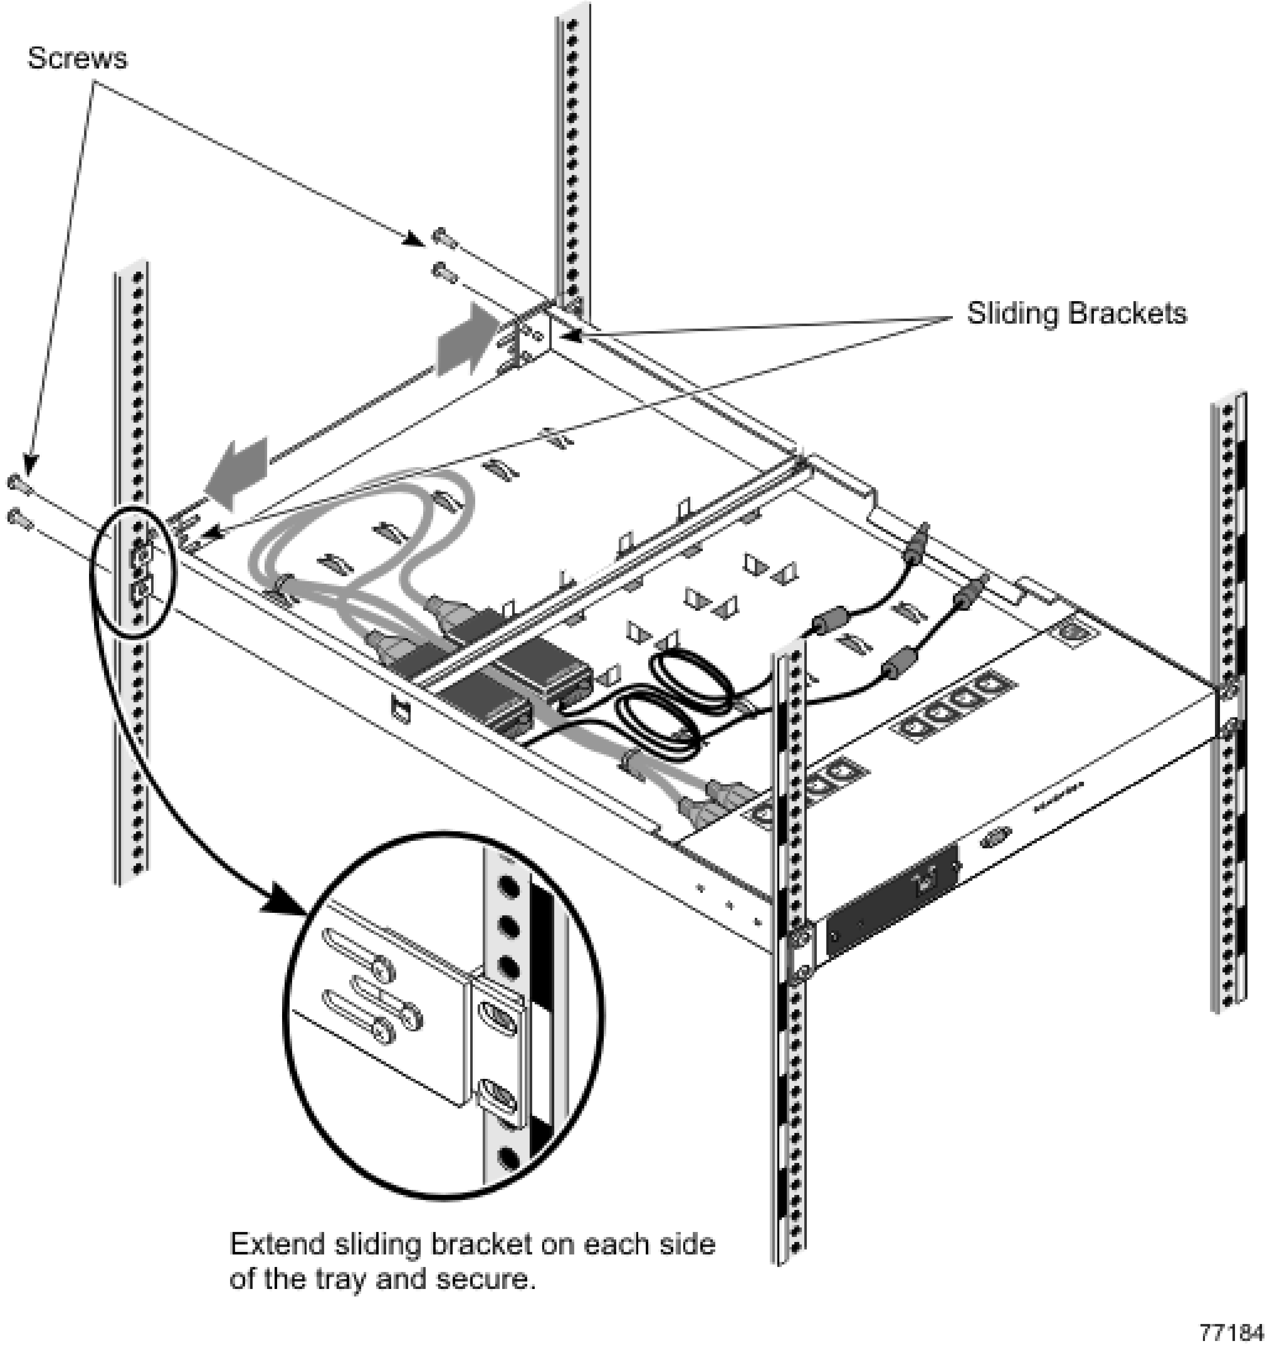

Figure 23. Install Power Management Module tray into cabinet from the rear of the cabinet, part 2

e. Tighten the mounting screws in the cabinet vertical rails with a Phillips #3 screwdriver

f. Tighten the sliding bracket screws

5. Attach cables and power up the Power Management Module.
   a. Plug the Power Management Module power cable into the PDU
   b. Plug the power cables from the power bricks on the tray into the Power Management Module
   c. Power up the port servers
   d. Plug in the remote power cables to the RSA cards
   e. Verify that all LEDs indicate that power is applied

      If power is not applied, use the **telnet** command to enable power to the Power Management Module. Then enable the devices through the menu system. The default settings are for outlets to be on when power is applied.
6. Configure and set up the Power Management Module

Configure and setup after device replacement

1. Configure the new Power Management Module with the same IP address of the Power Management Module you removed
2. Verify that the firmware level is correct.
3. IBM Customer Support personnel will update the firmware if necessary.
4. For firmware versions prior to 2.2, an SNMP patch must be applied to maintain security. The SNMP patch for the AP9212 is available for download at: http://apcc.com/tools/download/

Installation of power bricks for RSAs on tray

1. Place the power bricks on the tray, as shown in the figure below. Use the positioning tabs to place the power bricks. Make sure the power bricks are positioned so that the label is facing up, as shown in Figure 24 on page 118.

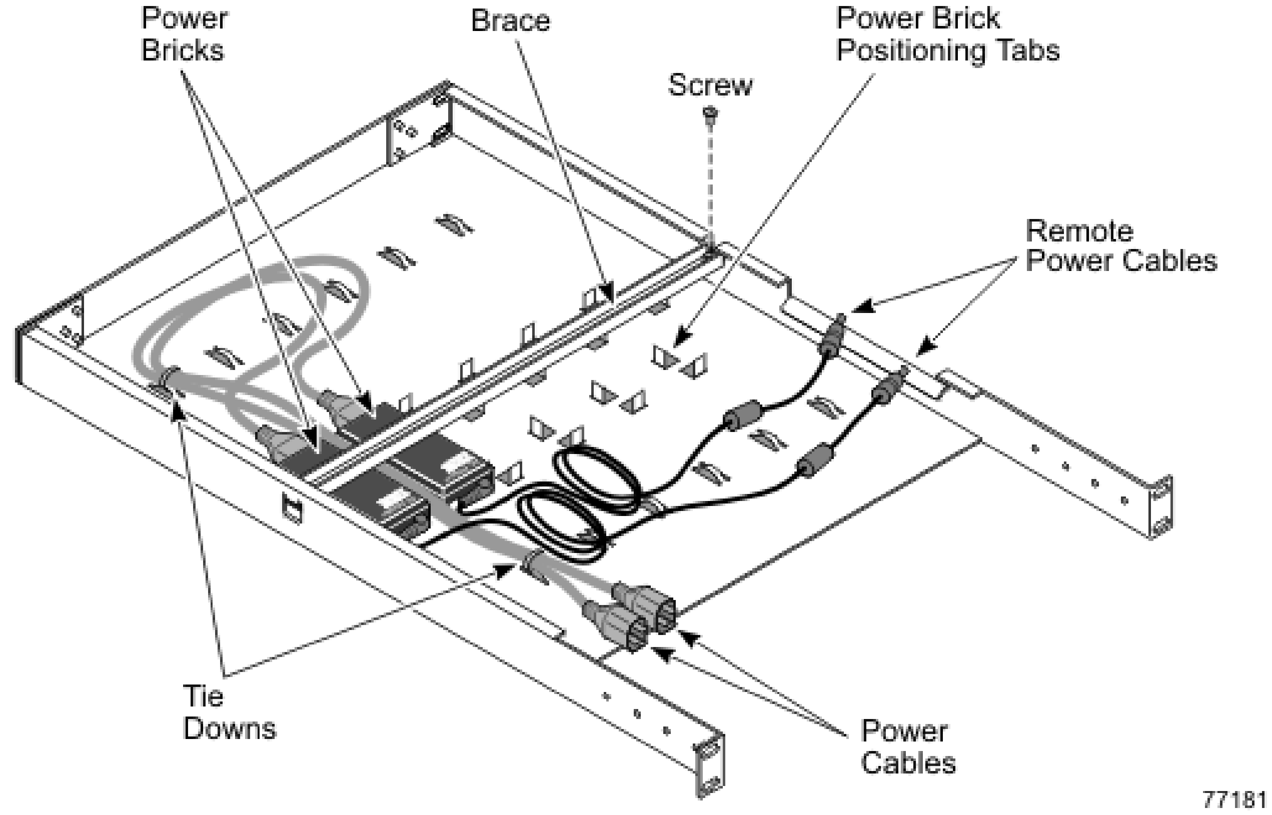

Figure 24. Mount power bricks onto Power Management Module tray

2. Mount the brace over the power bricks, as shown in Figure 21 on page 114. Mount it so that one end fits under the notch in the tray and one end fits over a screw hole. Install and tighten the screw to secure the brace.

Related topics

Additional product information: use the country selector for your locale or language
http://www.apcc.com/template/country_selection.cfm

The online search engine for APC documentation is available at:
http://sturgeon.apcc.com/techref.nsf/umanuals?openform

Use AP 9212 as a search parameter. Relevant documents include:

- *MasterSwitch Power Distribution Unit Installation and Quick Start Manual* (English/French/Spanish/Japanese/German)
- *MasterSwitch User's Guide* (English)

The online search engine for a variety APC technical questions including troubleshooting is available at:
http://www.apcc.com/support/index.cfm

Chapter 21. Power Distribution Unit removal and replacement

Contents

Removal and replacement

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

- Rack PDUs
- Front-end PDUs

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

To remove the Power Distribution Units take the following steps:

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

© Copyright IBM Corp. 2003

Part 4. Appendixes

© Copyright IBM Corp. 2003

Appendix A. Frequently Asked Questions

Contents

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

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

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

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

A: Sometimes the RSA cards get hung. They can be reset via the web interface, telnetting to the RSA card, or issuing the command ′**rpower -n <node>resetsp_hcp**.

Q: Why do I get the *′no hard drive found′* message from Linux when installing an x335 with IDE drives?

A: The file */csminstall/csm/1.3.0/kickstart.RedHat7.3/172.20.3.x-kickstart* needs to be changed.where *x* is the IP address of a node with IDE drives. Replace the references to *sda* with *hda*. Try the installation again.

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

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

Q: Why doesn't the x345 boot PXE correctly?

A: You cannot have a PCI ethernet card that uses the e1000 driver in the x345 when installing. Take the card out and retry the installation.

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

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

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

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

© Copyright IBM Corp. 2003

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

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

Q: What is causing Suse/Sles to continously install the nodes?

A: Check that fully qualified names (host.domainname) are used in the */etc/hosts* file and that the command **dnsdomainname** returns the correct domainname. Also make sure that */etc/dhcpd.conf* file contains the line: *'option domain-name "cluster.net";'* Once these changes have been made run the **csmsetupsis** 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. Now using the console watch the boot and the error messages should be on the console when the process has completed.

Q:Why do SuSE/SLES installs take forever?

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

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

A:Modify the */tftpboot/pxelinux.cfg/*.sis* file for the node you are trying to install and remove the *console=ttyS0,9600* line. Then use the KVM and switch to that node and you can see the error msg.

Appendix B. Error Logs

Contents

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

Linux Logs

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 Logs

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 bootup. 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
```

CSM Logs

CSM log files can be viewed in*/var/log/csm/installnode.log*

RSA Log

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

APC Event Log

You can view the 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

© Copyright IBM Corp. 2003

Appendix C. Known problems

Contents

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 to the x335 Service Processor Firmware diskette.
2. Select Configuration Settings at the Main Menu.
3. Select General Settings at the Configuration Menu.
4. Set the 75% Full and Log Full setting to **No**.

COM port settings in BIOS

The COM Port settings for the cluster node should be:

COM Port 1/A
2E8

COM Port 2/B
2F8

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

CSM

Stale NFS mounts

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

rpower hard shut down

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

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

Storage

Driver module ordering

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

© Copyright IBM Corp. 2003

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

Subsequent KVMs unresponsive

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

RSA and Service Processor

If there are any RSA errors, check to ensure the RSA is in PCI slot 2.

RSA unable to load firmware

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

RSA/Service Processor invalid naming

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

Light path points to PCI LED

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

Myrinet

Myrinet communication fails

If communication fails over Myrinet then check the following:

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

Appendix D. Setting up network switches

Contents

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 4006's.

When setting up VLANs on a 4006 running Catos make sure to set the vtp domain name. This can be any name since we are not using vtp to maintain the VLANs.

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.

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

4006 running Catos has one port that can be used for an Ethernet connection. It is sc0 and can be in any VLAN. For our purposes it should 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. It is the same story as above for the management port on the SUP card.

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.

RJ45 (copper) adapter GBICS must be connected to a gigabit port or it will not link. Those GBICS will not negotiate speed.

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

Switch commands

Switch commands for 3508/3550 running IOS

These commands will work with a 4006 running IOS as well. To set up VLANs issue the following commands:

© Copyright IBM Corp. 2003

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

To assign ports to the vlan issue the following commands:

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

To set Ethernet address for switch assign to Management VLAN issue the following commands:

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

To assign a name to the switch issue the following commands:

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

To see VLAN setup issue the following command:

```
show vlan
```

Other debug commands: To see spanning tree info on a port by port basis issue the following command:

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

Switch commands for 4006 running IOS

These commands will work with 3550 running IOS as well. To set up VLANs issue the following commands:

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

To assign ports to the VLAN issue the following commands:

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

To set Ethernet address for switch assigned to Management VLAN issue the following commands:

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

To assign a name to the switch issue the following commands:

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

To create an Etherchannel issue the following commands:

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

To remove an Etherchannel issue the following commands:

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

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

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

This command will generate an error message about port differences. Once the command works, set it back to the desirable mode.

To turn off the spanning tree on ports going to Cluster and Storage Nodes issue the following commands:

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

To see the VLAN setup issue the following command:

```
show vlan
```

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

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

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

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

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

```
show spanning-tree brief
```

See Etherchannels that are up and running:

```
show etherchannel
```

Switch commands for 4006 running CATOS

To set up VLANs issue the following commands:

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

To assign ports to the VLAN issue the following commands:

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

To set Ethernet address for switch assigned to Management VLAN issue the following command:

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

To set switch interface to a VLAN later issue the following command:

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

To create an Etherchannel issue the following command:

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

To assign a name to the switch

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

For the following command to work make sure all ports in the group are set up identically. If an Etherchannel doesn't form, debug using this command

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

This command will generate an error message about port differences. Once the command works, set it back to the desirable mode.

To turn off the spanning tree on ports going to compute and storage nodes issue the following commands:

```
set port host
```

To see the VLAN setup issue the following command:

```
show vlan
```

To set the switch as the spanning tree primary:

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

To set the switch as the spanning tree secondary:

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

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

```
show spantree
```

See Etherchannels that are up and running:

```
show channel
```

If an Etherchannel doesn't link up, first disable and then enable the ports:

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

Miscellaneous CISCO switch commands for CATOS

To clear configuration information from all modules in the switch issue the following command:

```
clear config <all>
```

To clear configuration information from a module issue the following command:

```
clear config <mod>
```

To see what ports are blocked by spanning tree issue the following command:

```
show spantree
```

Miscellaneous CISCO switch commands for IOS

To see what ports are blocked by the spanning tree issue the following command:

```
show sp br
```

Appendix E. International License Agreement for Non-Warranted Programs

Contents

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 ensure 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.

© Copyright IBM Corp. 2003

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 condiion 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 guaranty″ 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: Embedded Software from Cisco Systems, Inc.

Program-unique Terms

The following additional Software License terms supplement the IBM International License Agreement for Non-Warranted Programs (ILA) provided to you by IBM or an IBM reseller in connection with your purchase of an IBM product. Solely with respect to your use of the Cisco software (the ″Cisco Software″) contained within the IBM product you have purchased, these terms supersede the ILA.

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.

Notices

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

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

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

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

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

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

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

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

Trademarks

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

@server
IBM
IBMLink™
xSeries
Cluster 1350

© Copyright IBM Corp. 2003

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

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

Electronic emissions 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

声　明

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

Japanese Voluntary Control Council for Interference (VCCI) statement

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

Regulatory and Compliance Requirements

The IBM eServer Cluster 1350 systems meet the following regulatory and compliance requirements. Unless otherwise stated, the information applies to all systems.

Telecom

Three communication interfaces exist in the IBM eServer Cluster 1350. Each interface is an alternative configuration of the PCI-360 High Speed Serial Interface.

Only the SCC-P X.21 interface is used to support X.25 communication. This interface is certified to European X.25 requirements. The board is labeled appropriately for connection to Telecom networks throughout Europe. The shipping container is labeled with the CE label.

The interface is also certified for X.25 communication in Australia.

Safety Compliance

USA

IBM eServer Cluster 1350 systems have third-party certification to UL 60950, Safety of Information Technology Equipment.

These systems can include components such as an FC Host Adapter PCI card, the FC Switch, the FC-AL Hub, and the FC Bridge that contain laser component-assemblies called GLM (Gigabaud Link Module), GBIC (Gigabit Interface Converter), and GOT (Gigabit Optical Transceiver).

All models of laser component-assemblies comply with the requirements for Class-1 laser products set by the Department of Health and Human Services (DHHS) regulation 21 CFR Subchapter J. This compliance is indicated by markings on the laser component-assembly. Note that the Class 1 conformity label may not be visible when the laser component-assembly is installed in the system.

Canada

IBM eServer Cluster 1350 systems have third-party certification to CSA 22.2 #60950, Safety of Information Technology Equipment.

International

IBM eServer Cluster 1350 systems have third party certification, and the SQuad, MQuad, and Bootbay are self-certified to EN60950 and comply with IEC 60950, Safety of Information Technology Equipment.

These systems also comply with EN60825, Class 1; and with IEC 825, Class 1: Radiation Safety of Laser Products. Systems can include components such as an FC Host Adapter PCI card, the FC Switch, the FC-AL Hub, and the FC Bridge that contain laser component-assemblies called GLM (Gigabaud Link Module), GBIC (Gigabit Interface Converter), and GOT (Gigabit Optical Transceiver).

All models of laser component-assemblies are certified as Class-1 laser products that conform to the requirements contained in the International Electrotechnical Commission (IEC) standard 825 and CENELEC (European Committee for Electrotechnical Standardization) European Normalization standard EN 60825. All these assemblies are certified by the German testing institute VDE or an equivalent agency approved by the European Union.

Note that the Class 1 conformity mark may not be visible when the laser component-assembly is installed in the system.

Batteries

IBM eServer Cluster 1350 systems and products provided by IBM for connection to IBM products may contain sealed lead acid batteries or nickel-cadmium batteries. These batteries must be recycled or disposed of properly. Recycling facilities may not be available in your area. In the United States, IBM has established a collection process for reuse, recycling, or proper disposal of used sealed lead acid, nickel cadmium and nickel metal hydride batteries and battery packs from IBM equipment. For information on proper disposal of the batteries in this product, please contact IBM at 1-800-426-4333. For information on disposal of sealed lead acid or nickel cadmium batteries outside the United States, contact your local waste disposal or recycling facility.

Environmental Statement

IBM eServer Cluster 1350 systems were designed with customers' concerns about the environment in mind. The environmental impacts of product shipping, usage and recycling were considered throughout. Some of these features are noted below. IBM is continuously working with its partners to improve the environmental quality of OEM products; but at this time environmental achievements cannot be assumed for all OEM devices.

Energy Savings

Four-processor boards do the work of several of the fastest dual-processor boards in previous multiprocessor systems. With each new processor generation, IBM offers computing power increases while achieving power consumption decreases on a per-unit basis. An extra benefit of this performance is energy savings for the user. These systems also continue to provide the Energy Star option on all video display terminals.

Hazardous Materials

IBM has eliminated the following environmentally hazardous materials from IBM eServer Cluster 1350:

- CFCs and HCFCs
- PCB and PCT
- Mercury
- Cadmium in packaging and ink
- Lead in plastic parts that weigh more than 25 grams
- PBB, PBBO, PBBE and PBDE in plastic parts that weigh more than 25 grams

Printed Circuit Boards

IBM continually reviews printed circuit board processes for opportunities for improvement. Boards used in the IBM eServer Cluster 1350 systems use the following environmentally conscious processes:

- No solvent cleaners are used.
- Aqueous solutions are used for solder mask and photo imaging processes.
- No lead is used in the surface finish.
- Materials returned to IBM are reused and recycled such that only 5% of the base board volume, as nontoxic ash, goes to a landfill.

Documentation

On-line documentation is now available to all system users, thus minimizing the need for printed manuals. Manuals that are printed use at least 20% recycled, alkaline bleached paper and water based ink. The manuals are 95% recyclable.

Packaging

Shipping containers for IBM eServer Cluster 1350 systems and their components have been designed to meet the German Packaging Ordinance and the (US) Institute of Packaging Professionals Environmentally Responsible Packaging Handbook, R3P2. This compliance includes the following features:

- Wood is minimized, untreated and not permanently secured to unlike material.
- Cardboard and paper are free from foreign materials that impede recycling.
- Foams are CFC-free and are not permanently secured to unlike material.
- Plastics are free from foreign materials that impede recycling, with the exception of static bags.
- Commingled or copolymer plastics are not used, with the exception of static bags.
- The sum concentration of incidental levels of lead, cadmium, mercury and hexavalent chromium is less than 0.01% by weight.
- Inks are water based with no heavy metal additives.
- The total package is 100% recyclable with the exception of static bags and polyurethane foam, both of which are reusable.
- The total package is made of 10-20% recycled materials.

Upgradability

The modular design of the IBM eServer Cluster 1350 systems and their adherence to industry standards allow the systems to be both scalable and easily upgradable. Features include scalable memory, PCI I/O cards, standard 19 inch rack-mount capability, and clustering for processing units.

Recycling

IBM eServer Cluster 1350 systems cabinets and peripherals all have steel enclosures. The only significant plastic parts are the decorative cabinet doors and top panel. Waste material from their molding process is recycled into production parts. In addition, for ease of recycling, the ISO 11469 recycling mark designating the plastic composition is displayed on the doors and panels.

Index

© Copyright IBM Corp. 2003

Printed in U.S.A.