PENGUIN SOLUTIONS | **Stratus**
ztC Endurance™

ztC Endurance 3xx0, 5xx0, 7xx0, and 9xx0 Systems: User's Guide (Linux) (R007C)

Notice

The information contained in this document is subject to change without notice.

UNLESS EXPRESSLY SET FORTH IN A WRITTEN AGREEMENT SIGNED BY AN AUTHORIZED REPRESENTATIVE OF PENGUIN SOLUTIONS, PENGUIN SOLUTIONS MAKES NO WARRANTY OR REPRESENTATION OF ANY KIND WITH RESPECT TO THE INFORMATION CONTAINED HEREIN, INCLUDING WARRANTY OF MERCHANTABILITY AND FITNESS FOR A PURPOSE.

Penguin Solutions assumes no responsibility or obligation of any kind for any errors contained herein or in connection with the furnishing, performance, or use of this document. Software described in Penguin Solutions' documents (a) is the property of Penguin Solutions (PENG) Ireland Limited or the third party, (b) is furnished only under license, and (c) may be copied or used only as expressly permitted under the terms of the license.

Penguin Solutions' documentation describes all supported features of the user interfaces and the application programming interfaces (API) developed by Penguin Solutions. Any undocumented features of these interfaces are intended solely for use by Penguin Solutions personnel and are subject to change without warning.

Copyrights

Stratus, the Stratus logo, Stratus ztC, the Stratus Technologies logo, the Stratus 24x7 logo, Stratus ztC Endurance, Smart Exchange, and Automated Uptime Layer with Smart Exchange are trademarks of Stratus Technologies Ireland, Ltd.

Intel and the Intel Inside logo are registered trademarks and Xeon is a trademark of Intel Corporation or its subsidiaries in the United States and/or other countries/regions.

Microsoft, Windows, Windows Server, and Hyper-V are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries/regions.

VMware, vSphere, and ESXi are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions.

The registered trademark Linux is used pursuant to a sublicense from the Linux Mark Institute, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Google and the Google logo are registered trademarks of Google Inc., used with permission. The Chrome browser is a trademarks of Google Inc., used with permission.

Mozilla and Firefox are registered trademarks of the Mozilla Foundation.

Red Hat is a registered trademarks of Red Hat, Inc. in the United States and other countries.

Ubuntu and Canonical are registered trademarks of Canonical Ltd.

All other trademarks and registered trademarks are the property of their respective holders.

Table of Contents

ztC Endurance 3xx0, 5xx0, 7xx0, and 9xx0 Systems: User's Guide (Linux)

The ztC Endurance 3xx0, 5xx0, 7xx0, and 9xx0 Systems: User's Guide (Linux) (R007C) describes ztC Endurance systems, how to deploy them, and how to use them.

Part 1: ztC Endurance Hardware Guide describes the ztC Endurance hardware. It includes:

- "Overview of Stratus ztC Endurance Systems" on page 17

- "Planning a Site for ztC Endurance Systems" on page 37

- "Installing and Connecting the ztC Endurance System Hardware" on page 51

- "Operating the ztC Endurance System Hardware" on page 85

- "Removing and Replacing System Components" on page 101

Part 2: ztC Endurance Software Guide describes the ztC Endurance software. It includes:

- "System Software Installation and Upgrade (Linux)" on page 143

- "System Software Administration (Linux)" on page 196

- "Using the ztC Endurance Console" on page 217

- "Overview of the BMC Web Console" on page 271

Part 3: ztC Endurance Supporting Documents describes the ztC Endurance supporting documents. It includes:

- "Release Notes: Automated Uptime Layer with Smart Exchange Release 2.0.0.0 (Linux)" on page 279

- "System Reference Information" on page 289

- "System Specifications" on page 289

Next Steps:

To install the system hardware, see "Installing and Connecting the ztC Endurance System Hardware" on page 51.

To install the system software, see "Installing the System Software" on page 143.

Related Topics

- ztC Endurance 3xx0, 5xx0, 7xx0, and 9xx0 Systems: Deploying a System (R002C)

Part 1: ztC Endurance Hardware Guide

The ztC Endurance Hardware Guide discusses the following topics:

- "Overview of Stratus ztC Endurance Systems" on page 17

- "Planning a Site for ztC Endurance Systems" on page 37

- "Installing and Connecting the ztC Endurance System Hardware" on page 51

- "Operating the ztC Endurance System Hardware" on page 85

- "Removing and Replacing System Components" on page 101

- "Electrical Circuit and Wiring Information" on page 131

Related Topics

- "ztC Endurance Software Guide" on page 143

- "ztC Endurance Supporting Documents" on page 277

1

Chapter 1: Overview of Stratus ztC Endurance Systems

A Stratus ztC Endurance system consists of a system chassis that you install on rails in a cabinet. At the front and rear of the system chassis, you view LEDs and access modules for maintenance. Each module is a customer-replaceable unit (CRU) that you can remove and replace.

Read the following topics to familiarize yourself with your ztC Endurance system:

- "System Front" on page 17
- "System Rear" on page 18
- "System CRU Modules" on page 20
- "Peripheral Components" on page 35
- "Planning a Site for ztC Endurance Systems" on page 37

Next Steps

When you are familiar with the system, install the system hardware as described in "Installing and Connecting the ztC Endurance System Hardware" on page 51.

> **Note**: As you read through the installation topics, the sections labeled **Next Steps** provide links that help you move to the next logical step in a procedure.

System Front

The front of the system consists of two storage modules and two compute modules. Each storage module and each compute module is a CRU module. The system front also includes the control-panel ear on the right and a blank ear on the left. See Figure 1-1.

For more information, see:

- "System CRU Modules" on page 20 for information about the compute module (which contains RDIMMs), and the storage module (which contains disk drives)

- "Status LEDs" on page 90 for information about the LEDs on the control-panel ear, the storage module, and the compute module

- "System Power" on page 86 for information about the power button on the control-panel ear

- "Installing and Connecting the ztC Endurance System Hardware" on page 51 for information about installing and connecting the hardware

- "System Specifications" on page 289 for detailed specifications about the system and its components

Figure 1-1: System Front

1	Compute module A	4	Storage module B
2	Storage module A	5	Control-panel ear
3	Compute module B		
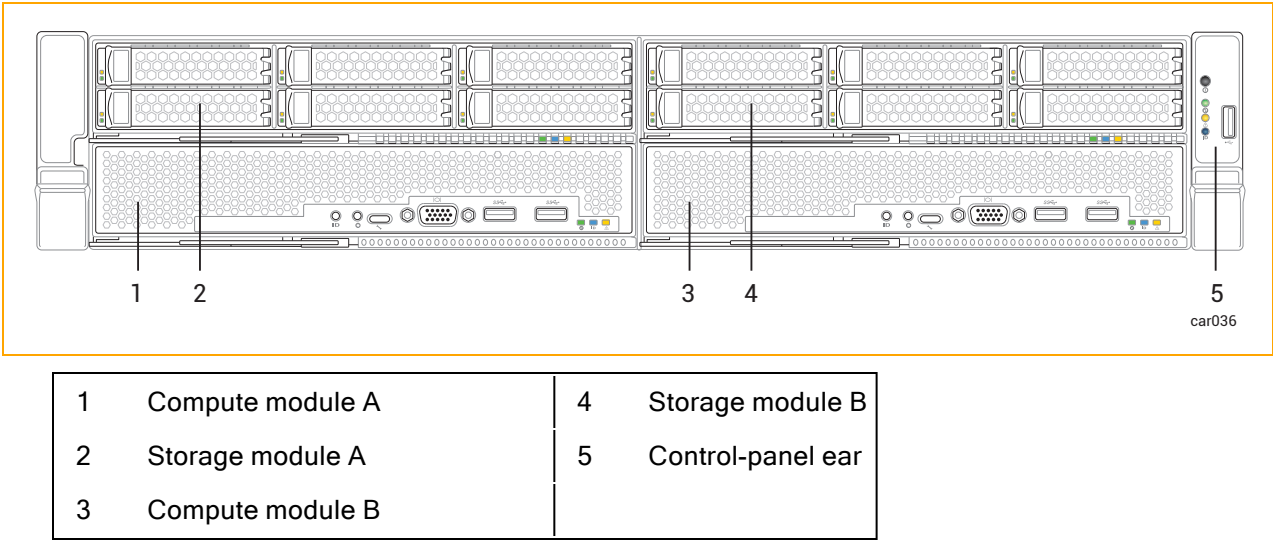
> **Note**: The left ear contains a QR code that links to ztC Endurance Release Notes and documentation at https://ztcendurancedoc.stratus.com.

Related Topics

- "Planning a Site for ztC Endurance Systems" on page 37

- "System Reference Information" on page 289

System Rear

The rear of the system consists of two power-supply units (PSUs) and two I/O modules. Each PSU and each I/O module is a CRU module. See Figure 1-2.

For more information, see:

- "System CRU Modules" on page 20 for information about the I/O modules and PSUs

- "Status LEDs" on page 90 for information about the LEDs on the I/O module

- "Installing and Connecting the ztC Endurance System Hardware" on page 51 for information about installing and connecting the hardware

- "System Specifications" on page 289 for detailed specifications about the system and its components

Figure 1-2: System Rear
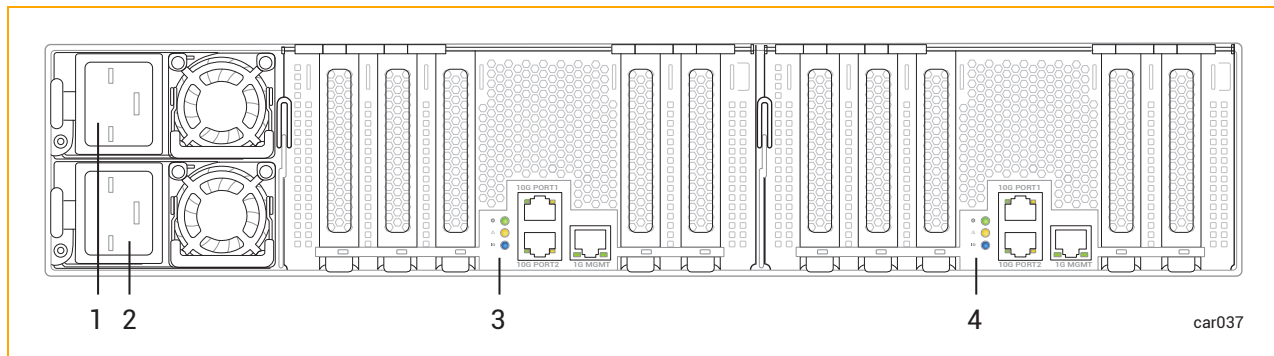
1	Power-supply unit (PSU) A	3	I/O module A
2	Power-supply unit (PSU) B	4	I/O module B

Related Topics

- "Planning a Site for ztC Endurance Systems" on page 37

- "System Reference Information" on page 289

System Ports

Ports are located at the front (Figure 1-3) and rear (Figure 1-4) of the system.

> **Note**: The USB 2.0 port on the control-panel ear is the only front port that stays with the active compute module through a Smart Exchange operation. That is, after the active compute module becomes the standby module, the USB 2.0 port automatically connects to the new active compute module.

Figure 1-3: System Ports: Front
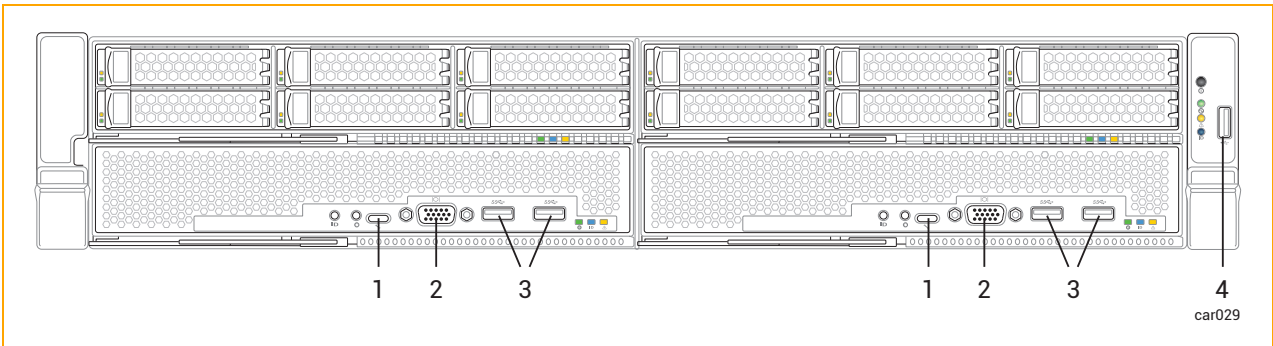
1	Debug ports (2)	3	USB 3.0 ports (4)
2	VGA (monitor port) (2)	4	USB 2.0 port (1)

Figure 1-4: System Ports: Rear
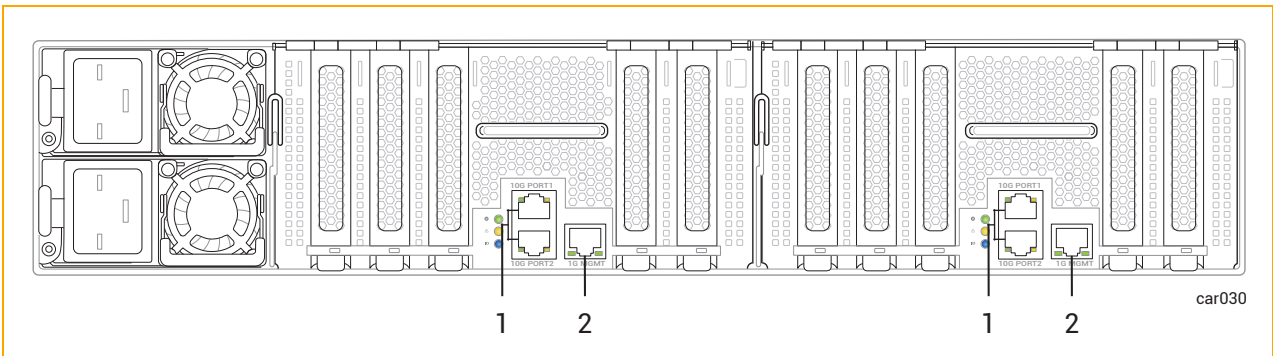
1	10G ports (4)
2	1G MGMT ports (2)

Related Topics

- "System Front" on page 17

- "System Rear" on page 18

System CRU Modules

A ztC Endurance system contains customer-replaceable unit (CRU) modules, which allow on-site personnel to easily remove and replace faulty components with minimal training or tools. These CRU modules are:

- Compute module and its RDIMMs

- Storage module and its disk drives

- I/O module and its PCIe adapters

- PSUs

Related Topics

- "System Front" on page 17

- "System Rear" on page 18

- "Removing and Replacing System Components" on page 101

- "System Specifications" on page 289

Compute Module

A ztC Endurance system contains two compute modules: compute module A and compute module B. Figure 1-1 shows the location of each compute module in a system.

A compute module is a 2U component with processors, RDIMMs, BIOS, BMC, fans, and other technology to boot and run the ztC Endurance system.

Figure 1-5 shows an individual compute module. Each compute module provides:

- Ports, including one VGA port and two USB 3.0 ports. See "System Ports" on page 19. (The UART USB-C connector, known as the debug port, is used only by Penguin Solutions Customer Service for debugging.)

- An ID button, which activates the ID LED on the control-panel ear. See "Status LEDs" on page 90.

- An NMI button. See "Using the NMI Button" on page 100.

- Status LEDs, which indicate the status of the compute module. See "Status LEDs" on page 90.

- RDIMMs, which provide the system's random-access memory (RAM). See "Replacing and Adding RDIMMs" on page 109.

Figure 1-5: Compute Module
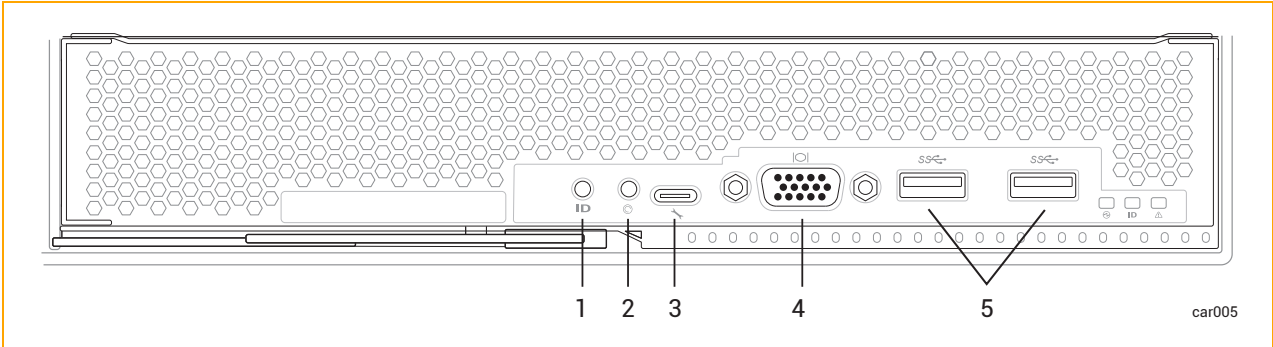
1	ID button	4	VGA port
2	NMI button	5	USB 3.0 port (x2)
3	Debug port		

Related Topics

- "System Front" on page 17

- "Planning a Site for ztC Endurance Systems" on page 37

- "Removing and Replacing a Compute Module" on page 106

- "System Specifications" on page 289

Memory in ztC Endurance Systems

Read this topic before you add or replace registered dual in-line memory modules (RDIMMs) in a ztC Endurance system.

Memory Components in Each Compute Module

Table 1-1 lists the memory components in each compute module.

Table 1-1: Memory Components

Model	Processors	Sockets	RDIMM Slots	Minimum Number of RDIMMs	Maximum Number of RDIMMs
31x0	1	2	16[1]	1	4
51x0	2	2	16	2	8
71x0	2	2	16	4	16
9110	2	2	16	16	16

Figure 1-6 shows the area of a compute module that contains memory components. Specifically, the figure shows the two processors, each installed in a socket, and the location of each RDIMM slot.

Figure 1-6: Location of Memory Components in a Compute Module
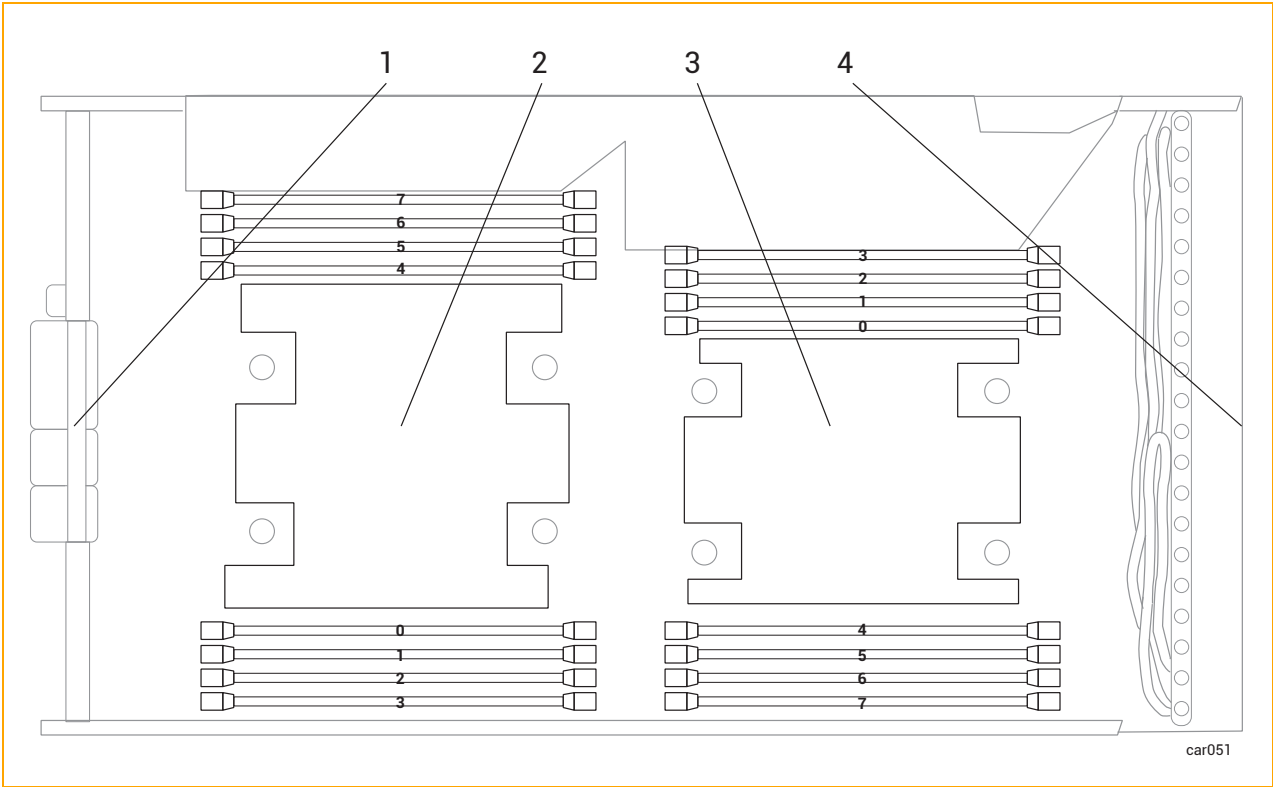
1	Rear of compute module	3	Processor 1
2	Processor 0	4	Front of compute module

Memory in ztC Endurance 31x0 Systems

Table 1-2 shows a representation, starting from the front of the compute module, of RDIMM locations for supported ztC Endurance 31x0 configurations. An "X" indicates the presence of an RDIMM; all other RDIMM slots are empty.

> **Note**: Table 1-2 does not show the RDIMM slots that are next to the empty socket.

Table 1-2: Memory in ztC Endurance 31x0 Systems

	ZEN31x0-64	ZEN31x0-128	ZEN31x0-256
Slot 7			
Slot 6		X	X
Slot 5			
Slot 4			X
Processor 0			
Slot 0	X	X	X
Slot 1			
Slot 2			X
Slot 3			

Memory in ztC Endurance 51x0 Systems

Table 1-3 shows a representation, starting from the front of the compute module, of RDIMM locations for supported ztC Endurance 51x0 configurations. An "X" indicates the presence of an RDIMM; all other RDIMM slots are empty.

Table 1-3: Memory in ztC Endurance 51x0 Systems

		ZEN51x0-128		ZEN51x0-256		ZEN51x0-512
Slot 7	Slot 3					
Slot 6	Slot 2			X		X	X
Slot 5	Slot 1					
Slot 4	Slot 0		X		X	X	X
Processor 0	**Processor 1**					
Slot 0	Slot 4	X		X		X	X
Slot 1	Slot 5					

		ZEN51x0-128		ZEN51x0-256		ZEN51x0-512	
Slot 2	Slot 6				X	X	X
Slot 3	Slot 7						

Memory in ztC Endurance 71x0 Systems

Table 1-4 shows a representation, starting from the front of the compute module, of RDIMM locations for supported ztC Endurance 71x0 configurations. An "X" indicates the presence of an RDIMM.

Table 1-4: Memory in ztC Endurance 71x0 Systems

		ZEN71x0-256		ZEN71x0-512		ZEN71x0-1024	
Slot 7	Slot 3					X	X
Slot 6	Slot 2	X		X	X	X	X
Slot 5	Slot 1					X	X
Slot 4	Slot 0		X	X	X	X	X
Processor 0	**Processor 1**					X	X
Slot 0	Slot 4	X		X	X	X	X
Slot 1	Slot 5					X	X
Slot 2	Slot 6		X	X	X	X	X
Slot 3	Slot 7					X	X

Memory in ztC Endurance 9110 Systems

Table 1-5 shows a representation, starting from the front of the compute module, of RDIMM locations for supported ztC Endurance 9110 configurations. An "X" indicates the presence of an RDIMM.

Table 1-5: Memory in ztC Endurance 9110 Systems

		ZEN9110-1024	
Slot 7	Slot 3	X	X

		ZEN9110-1024	
Slot 6	Slot 2	X	X
Slot 5	Slot 1	X	X
Slot 4	Slot 0	X	X
Processor 0	**Processor 1**	X	X
Slot 0	Slot 4	X	X
Slot 1	Slot 5	X	X
Slot 2	Slot 6	X	X
Slot 3	Slot 7	X	X

Related Topics

- "Replacing and Adding RDIMMs" on page 109

- "Compute Module" on page 21

- "Hardware Page" on page 229

[1]In a ztC Endurance 31x0 system, the eight RDIMM slots next to the empty socket cannot contain RDIMMs.

Storage Module

A ztC Endurance system contains two storage modules: storage module A and storage module B. Figure 1-1 shows the location of each storage module in a system.

Figure 1-7 shows an individual storage module. Each storage module provides:

- Six slots for disk drives. For additional information about disk drives, see:

    - "System Specifications" on page 289 for information about the supported disk drives.

    - "Disk Drives" on page 27 for information about disk-drive configuration and numbering.

    - "Internal Disk Drive Problems and Solutions" on page 99 for troubleshooting information.

    - "Removing and Inserting a Disk Drive" on page 113 for maintenance information.

- Status LEDs, which indicate the status of the storage module. See "Status LEDs" on page 90.

Figure 1-7: Storage Module
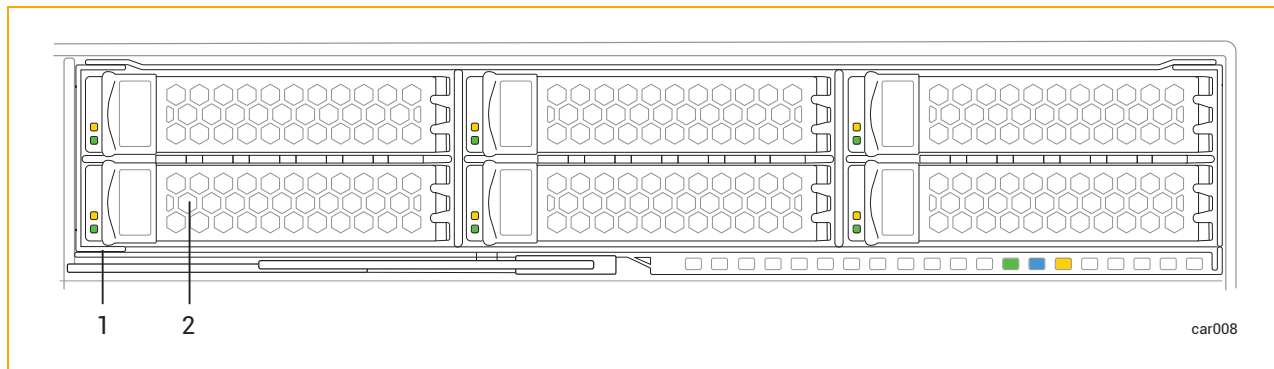
1	Storage module
2	Disk drive slot with disk drive (x6)

Related Topics

- "System Front" on page 17

- "Planning a Site for ztC Endurance Systems" on page 37

- "Removing and Replacing a Storage Module" on page 111

Disk Drives

When you receive your system from the manufacturer, disk drives that were ordered at the same time the system was ordered are already installed. To install additional disk drives, be sure to read "Installing Additional Disk Drives" on page 28 and "Important Considerations When Installing Disk Drives" on page 29 before you read "Removing and Inserting a Disk Drive" on page 113.

See "System Specifications" on page 289 for information about the supported disk drives.

Understanding Disk-Drive Numbering

Storage module A and storage module B each contain six disk-drive slots, for a total of 12 per system. The slots are labeled 1 through 6 on each storage module, as shown in Figure 1-8.

Figure 1-8: Disk-Drive Slot Numbering
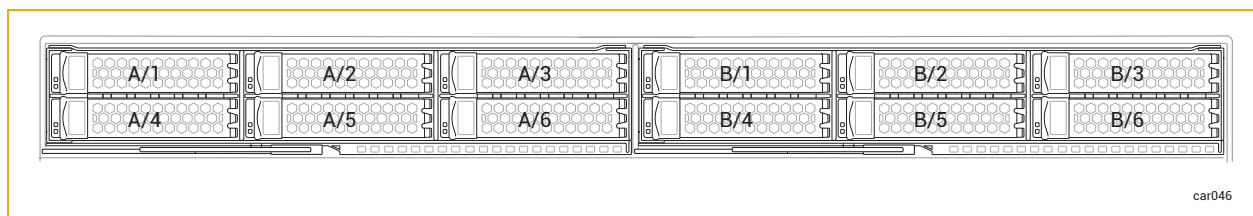
You can view information about the disk drives from "Hardware Page" on page 229 of the ztC Endurance console and from your operating-system software. The physical disk-drive slot numbering corresponds to the software numbering. For example, disk-drive slot 1 corresponds to slots A/1 and B/1 in the software, disk-drive slot 2 corresponds to slots A/2 and B/2, and so on.

Installing Additional Disk Drives

Install additional disk drives in your system in the order shown in Table 1-6. For example, if you are adding two disks to a system that already contains disks in slots A/1 and B/1, install the two additional disks in slots A/2 and B/2, for a total of four disks in the system.

Table 1-6: Disk-Drive Installation Order

Number of Disks	Storage Module A	Storage Module B
2	Slot 1	Slot 1
4	Slots 1, 2	Slots 1, 2
6	Slots 1, 2, 3	Slots 1, 2, 3
8	Slots 1, 2, 3, 4	Slots 1, 2, 3, 4
10	Slots 1, 2, 3, 4, 5	Slots 1, 2, 3, 4, 5
12	Slots 1, 2, 3, 4, 5, 6	Slots 1, 2, 3, 4, 5, 6

Important Considerations When Installing Disk Drives

Before you install disk drives, consider the following:

- The boot drive must be installed in slot 1.

- A storage module can contain different disk-drive models. However, you must install identical disk-drive models in matching storage module A and B slots. For instance, if you install a 1.6 TB disk drive in slot 1 of storage module A, you must install a 1.6 TB disk drive in slot 1 of storage module B.

- Make sure all disk drives are correctly inserted as described in "To insert a disk drive" on page 114.

Related Topics

- "Removing and Inserting a Disk Drive" on page 113

- "Internal Disk Drive Problems and Solutions" on page 99

- "Storage Module" on page 26

I/O Module

A ztC Endurance system contains two I/O modules: I/O module A and I/O module B. Figure 1-2 shows the location of each I/O module in a system.

Figure 1-9 shows an individual I/O module. Each I/O module provides:

- Ethernet ports. See "System Ports" on page 19.

  - One 1G (RJ45) port for an out-of-band management network

  - Two 10G (RJ45) ports

- Status LEDs, which indicate the status of the I/O module. See "Status LEDs" on page 90.

- Five PCIe 4.0 slots for half-height PCIe adapters. For additional information about PCIe adapters, see:

  - "System Specifications" on page 289 for information on the supported PCIe adapters.

  - "Communications Line Planning" on page 47 for information on the communications lines.

  - "PCIe Adapters" on page 30 for information about PCIe-adapter configuration and numbering.

  - "Removing and Installing a PCIe Adapter" on page 121 for maintenance information.

Figure 1-9: I/O Module

car006
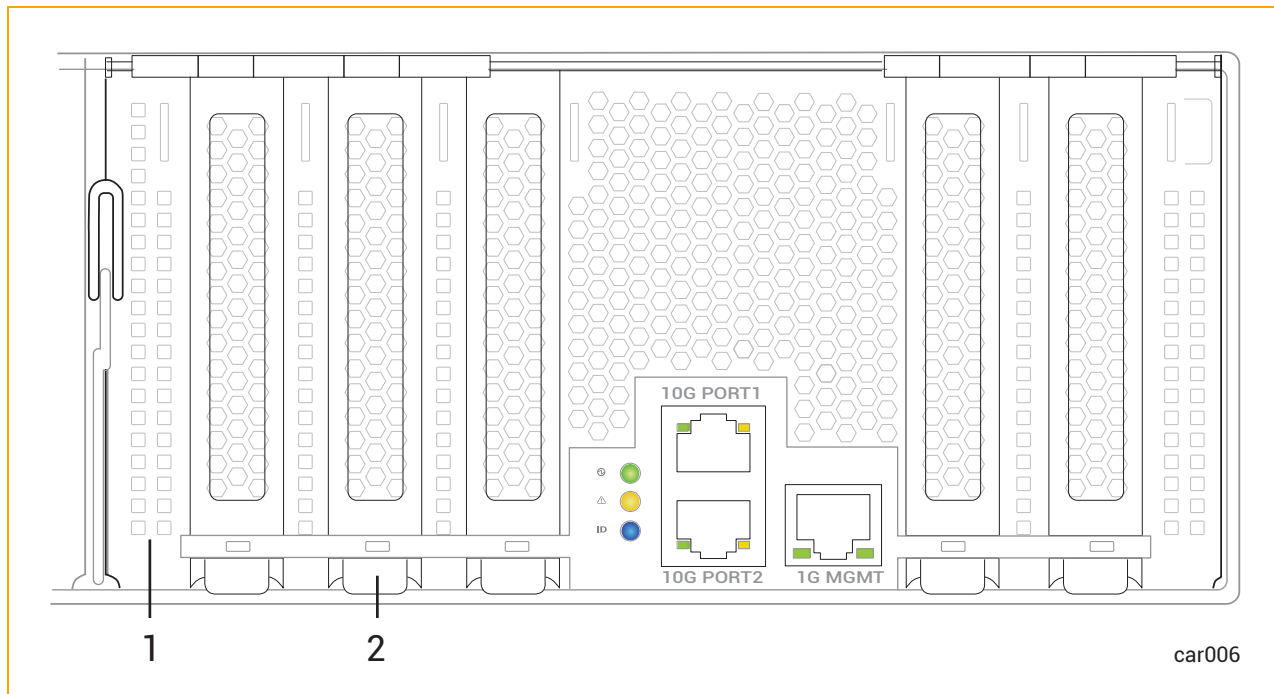
1	I/O module
2	PCIe adapter slot (x5)

Related Topics

- "System Rear" on page 18

- "Planning a Site for ztC Endurance Systems" on page 37

- "Removing and Replacing an I/O Module" on page 118

- "Hardware Page" on page 229

- "System Reference Information" on page 289

- "System Specifications" on page 289

PCIe Adapters

When you receive your system from the manufacturer, PCIe adapters that were ordered at the same time the system was ordered arrive pre-installed. To install additional PCIe adapters, be sure to read the following before you read "Removing and Installing a PCIe Adapter" on page 121

- "Understanding PCIe-Slot Numbering" on page 31

- "Installing Additional PCIe Adapters" on page 32

-

-

See for information about the supported PCIe adapters.

Understanding PCIe-Slot Numbering

I/O module A and I/O module B each contain five PCIe adapter slots, for a total of ten per system. The slots are labeled 1 through 5 on each I/O module, as shown in Figure 1-10. Each I/O module also contains an embedded dual-port 10G adapter.

Figure 1-10: PCIe-Adapter Physical Slot Numbering
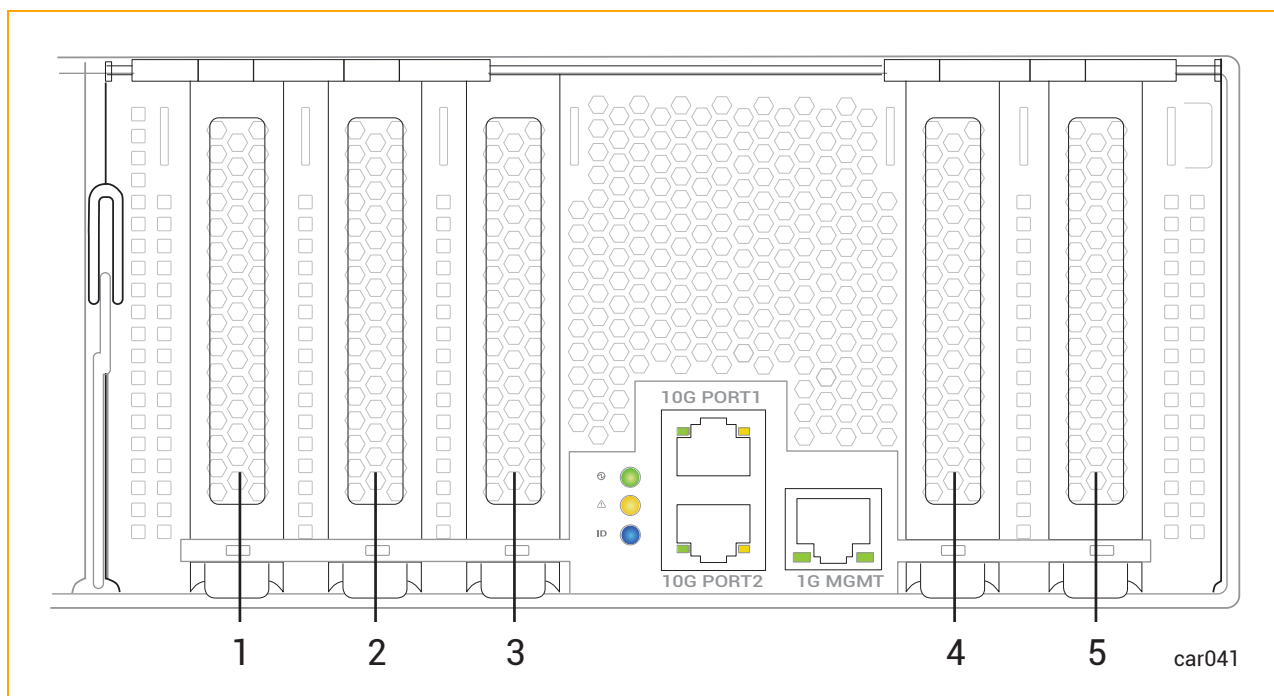
1	PCIe slot 1	4	PCIe slot 4
2	PCIe slot 2	5	PCIe slot 5
3	PCIe slot 3		

You can view information about the PCIe adapters from of the ztC Endurance console and from your operating-system software.

- For PCIe slots 1 through 5, the physical PCIe slot numbering corresponds to the software numbering. For example, PCIe slot 1 corresponds to slots A/1 and B/1 in the software, PCIe slot

2 corresponds to slots A/2 and B/2, and so on.

- The embedded dual-port 10G adapters appear in the software as slots A/6 and B/6.

Installing Additional PCIe Adapters

Install additional PCIe adapters in your system in the order shown in Table 1-7. For example, if you are adding two PCIe adapters to a system that already contains PCIe adapters in slots A/1 and B/1, install the two additional PCIe adapters in slots A/2 and B/2, for a total of four PCIe adapters in the system.

Table 1-7: PCIe-Adapter Installation Order

Number of Adapters	I/O Module A	I/O Module B
2	Slot 1	Slot 1
4	Slots 1, 2	Slots 1, 2
6	Slots 1, 2, 3	Slots 1, 2, 3
8	Slots 1, 2, 3, 4	Slots 1, 2, 3, 4
10	Slots 1, 2, 3, 4, 5	Slots 1, 2, 3, 4, 5

Important Considerations When Installing PCIe Adapters

Before you install PCIe adapters, consider the following:

- Use only PCIe adapters approved by Stratus. Before approving a PCIe adapter, Stratus thoroughly tests it to make sure that it interacts properly with Stratus systems, including I/O redundancy with a duplexed partner if an active PCIe adapter or I/O module fails or is removed from service. Customers who have attempted to use inadequately tested PCIe adapters in their systems have frequently experienced problems that proved difficult to isolate and diagnose and resulted in delayed projects.

- Thoroughly test any new PCIe adapter, including adapters sold or approved by Stratus, with your applications in a test environment before using it in a production environment.

- An I/O module can contain different PCIe adapter models. However, you must install identical PCIe adapter models in matching I/O module A and B slots. For instance, if you install an Intel X550-T2 PCIe adapter in slot 1 of I/O module A, you must install an Intel X550-T2 PCIe adapter in slot 1 of I/O module B.

Valid PCIe Configurations

Table 1-8 describes valid PCIe adapter configurations for each ztC Endurance I/O module. To interpret the table, read it from left to right.

> **Note**: Contact your authorized Penguin Solutions service representative if you have questions about the models supported by your ztC Endurance system.

Table 1-8: Valid PCIe Adapter Configurations for RHEL ztC Endurance Systems

ztC Endurance Model	Maximum Number of Each Adapter Type Allowed Per I/O Module				Maximum Number of Total Adapters Allowed Per I/O Module
	E810 Dual 100 Gbps	XXV710 Dual 25 Gbps	QLE2770 Single 32 Gbps	X550 or X710 Dual 10 Gbps	
Low-voltage 31x0	Not supported	Not supported	0	3	3
			1	1	2
Low-voltage 51x0	Not supported	Not supported	0	3	3
			1	1	2
31x0	Not supported	Not supported	1	3	3
51x0	1	2	2	4	4
71x0	1	2	2	5	5
9110	1	2	2	5	5

Related Topics

- "Removing and Installing a PCIe Adapter" on page 121
- "I/O Module" on page 29
- "Hardware Page" on page 229
- "System Specifications" on page 289

PSUs

A ztC Endurance system contains two power supply units (PSUs): PSU A and PSU B. Figure 1-2 shows the location of each PSU in a system.

Figure 1-11 shows an individual PSU. Each PSU includes a status LED, which indicates the status of the unit. See "Status LEDs" on page 90.

Figure 1-11: Power Supply Unit (PSU)
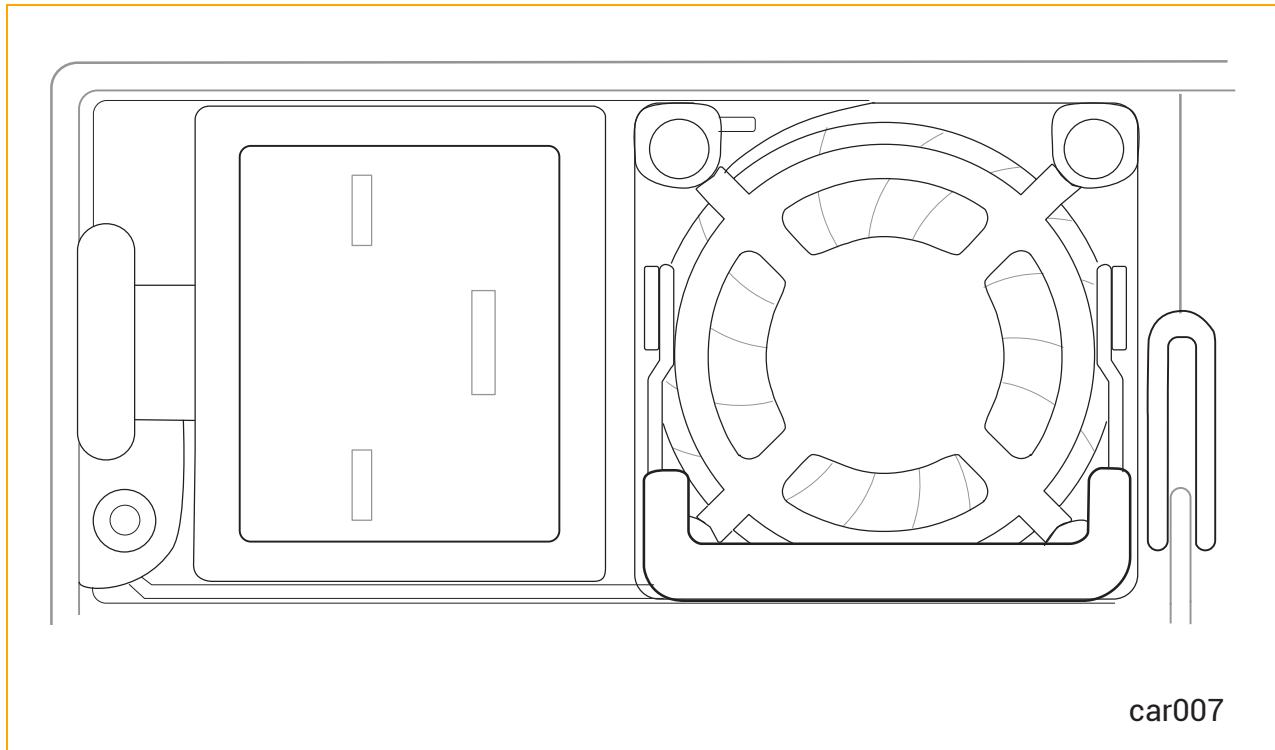
car007

Related Topics

- "System Rear" on page 18

- "Planning a Site for ztC Endurance Systems" on page 37

- "Removing and Replacing a PSU and Power Cords" on page 125

- "System Power" on page 86

- "Standby Power" on page 85

- "System Reference Information" on page 289

- "System Specifications" on page 289

Peripheral Components

Stratus supplies a mouse and keyboard for your ztC Endurance system.

For site-planning information for the mouse and keyboard, see "Site Planning for Peripheral Components" on page 49.

For information about connecting the mouse and keyboard, see Connecting Peripherals.

2

Chapter 2: Planning a Site for ztC Endurance Systems

> **Notes**:
>
> 1. See "Safety Notices" on page 52 for a list of important safety notices that you should review as part of your site planning.
>
> 2. All regulatory notices are provided in *Compliance Information for Stratus Products (R002G)*, which is available on StrataDOC at https://stratadoc.stratus.com/compliance_info/Compliance_Information_for_Stratus_Products.htm

Site planning for ztC Endurance systems includes:

- "Site Planning Checklist" on page 38

  Reviewing the site-planning checklist for a summary of questions you might need to consider when planning for the installation of a ztC Endurance system.

- "Electrical Power Planning" on page 40

  Providing electrical power sources that meet the requirements of the system and optional components, including the purchase of user-supplied power-distribution units (PDUs) and uninterruptible power supplies (UPSes).

- "Space Planning" on page 45

  Providing adequate space for the system or cabinet and for a desk or table to accommodate components outside a cabinet. Also providing enough space for servicing the systems and components.

Providing an environment that meets the system's requirements for ambient temperature and air quality.

- "Communications Line Planning" on page 47

Providing sufficient network lines and planning the location of the system and external components to accommodate the lengths of the connecting cables.

- "Cabinet Requirements" on page 47

Providing a cabinet that meets the system's requirements and provides sufficient space for all components, including user-supplied PDUs and UPSes.

- "Site Planning for Peripheral Components" on page 49

Providing a monitor, keyboard, and mouse (if applicable) that meet the system's requirements.

Providing for the electrical, cooling, power, and space needs of any peripheral components.

For additional documentation related to the ztC Endurance system, see the StrataDOC (ztC Endurance Version) page at https://ztcendurancedoc.stratus.com.

During the site planning and preparation processes, work closely with your facilities group or contractor to determine space, power, and environmental requirements. Enlist their help to provide a suitable location with sufficient alternating current (AC) power, heating, ventilation, and air conditioning (HVAC) capabilities, and network and telephone connections.

If the system is covered by a service agreement and you need help with site planning, contact your authorized Penguin Solutions service representative. If you have a contract with your authorized Penguin Solutions service representative to install the system, contact them after you have prepared the installation site and moved the system to the site. For more information about your authorized Penguin Solutions service representative, see "Getting Help" on page 295 or the **Penguin Solutions Customer Service** page at https://www.penguinsolutions.com/en-us/support.

See "System Specifications" on page 289 for the specifications of the ztC Endurance systems.

Site Planning Checklist

Referring to the information in this document, answer the following questions:

Planning for ActiveService Network Connectivity

☐ Will the ztC Endurance system connect to the Stratus ActiveService Network (ASN), allowing the system to automatically report problems to your authorized Penguin Solutions service representative?

☐ To allow connections to the ASN to report system health and events on your system, can you connect at least one pair of partnered Ethernet ports, one from each I/O module in the system, to a network that has Internet access?

☐ To optionally enable Penguin Solutions Customer Service or your authorized Penguin Solutions service representative to access the system to remotely diagnose, troubleshoot, and resolve problems online, can you connect the Ethernet port of the Baseboard Management Controller (BMC) in each module to a network that has Internet access? (The BMCs can communicate with the ASN even when the host operating system is not running or is unresponsive.)

Planning for Network Connectivity

☐ Each ztC Endurance system contains four embedded 10-gigabit Ethernet (GbE) ports, two in each I/O module, for connection to business networks. In addition, each system contains two embedded 1GbE ports that are dedicated to the BMCs for out-of-band management. How many embedded Ethernet ports will you connect to a network, and how many optional, non-embedded Ethernet PCIe adapters will you need to add to provide additional Ethernet bandwidth?

☐ How many embedded or non-embedded Ethernet ports in the system will require access to the Internet?

Planning AC Power

☐ Will you provide power to both sides of the system through a pair of PDUs?

☐ Will you protect the system power with a UPS?

☐ What optional components will you use?

☐ What are the AC power requirements of the system, including all optional components?

☐ What are the lengths and types of the power cords that are provided for the PDUs or system, and optional components?

☐ What type of AC receptacles do you need to provide?

☐ Is the AC power service wired properly?

Planning Space for the System

☐ If applicable, provide a table or desk for the monitor, keyboard, and mouse.

☐ Will the system and any external components fit where you plan to place them?

☐ What is the height of the cabinet you will use, and what is the total height of the systems and components that will be installed in the cabinet? Will the items fit into the cabinet? What components will be located outside the cabinet?

☐ What are the lengths and types of the interface and communications cables that will connect to the system?

☐ Have you created a sketch of how you plan to arrange the system at the installation site? Consider the available cable lengths, the placement of external devices, and the location of network and voice communication connections.

On the sketch, show the following:

- Location of the system and its external components

- Power cords, and telephone and interface cables

- Locations of AC power receptacles, Ethernet jacks, switches, and hubs

> **Note**: Make sure that all cords and cables are long enough to reach between their respective components and connectors. Route all cables out of the way of foot traffic.

Working with Other Groups

☐ Have you provided your facilities group and contractors with the sketch of how you plan to arrange the system and copies of the following?

- "Redundant Power Sources" on page 41

- "AC Power and HVAC Service Requirements" on page 41

- "Electrical Circuit and Wiring Information" on page 131

- Any notes you have about site planning

☐ Have you reviewed and discussed the requirements with the facilities personnel and contractors to ensure that all site modifications are understood and implemented?

If you have any questions about the number and types of components, contact your Penguin Solutions account representative or distributor.

Electrical Power Planning

For information about planning appropriate AC electrical power for the system and its peripheral components, see:

- "Redundant Power Sources" on page 41

- "AC Power and HVAC Service Requirements" on page 41

- "Electrical Power Connections" on page 42

-

-

Redundant Power Sources

ztC Endurance systems require at least two separate and independent AC power sources—an *A-side* power source and a *B-side* power source—that provide power to the system's power receptacles. Either source must be capable of continuing to provide power if power to the other source is lost.

The A-side power source provides power to the top power-supply unit (PSU) in the system. The B-side power source provides power to the bottom PSU.

Each PSU in the system is capable of providing power to the entire system. When both PSUs have electrical power, they share the load of the system components; however, if one PSU fails or if power is lost to either the A- or B-side power source, the remaining PSU can pick up the full load of the system.

Although either PSU is capable of powering the ztC Endurance system, it is still important to separate power sources as much as possible to maximize redundancy. Furthermore, if you connect user-supplied peripheral components that also contain two PSUs, attach each side to the independent A- and B-side power sources. For components that do not require two sources of power, such as monitors, management PCs, or other user-supplied components, you typically attach these components to the A-side power source.

If provided by the customer, a pair of PDUs may supply power to multiple systems and components in a cabinet. The top PDU in the cabinet typically provides A-side power. The bottom PDU typically provides B-side power.

AC Power and HVAC Service Requirements

Use the specifications in the "AC Power" section of "System Specifications" on page 289 to determine AC power and HVAC requirements for the ztC Endurance system. Because the ztC Endurance PSUs balance the load of the system components, the total power requirement for either side of the system should not exceed the system specifications listed in the "AC Power" section.

For any optional components, see the included documentation to obtain the nominal input line voltage (volts AC) and frequency (Hz) required.

> **Notice**: A branch circuit breaker with ground fault protection must allow a minimum of 3.5 milliamperes (mA) leakage current for each power cord.

See "System Specifications" on page 289 for additional ztC Endurance system specifications.

Electrical Power Connections

Provide two AC power sources that are as electrically independent of each other as the installation site allows. At a minimum, the two power sources must be powered by separate circuit breakers (maximum of 20A) to AC power and, if possible, be independent of each other beyond that level. The more electrical separation between the two power sources, the less likely they will both fail at the same time. Due to redundancy in ztC Endurance systems, power to either side of the system keeps the system in operation.

Figure 2-1 shows an example of how to connect a ztC Endurance system directly to two separate AC power sources.

Figure 2-1: Connecting a System Directly to AC Power
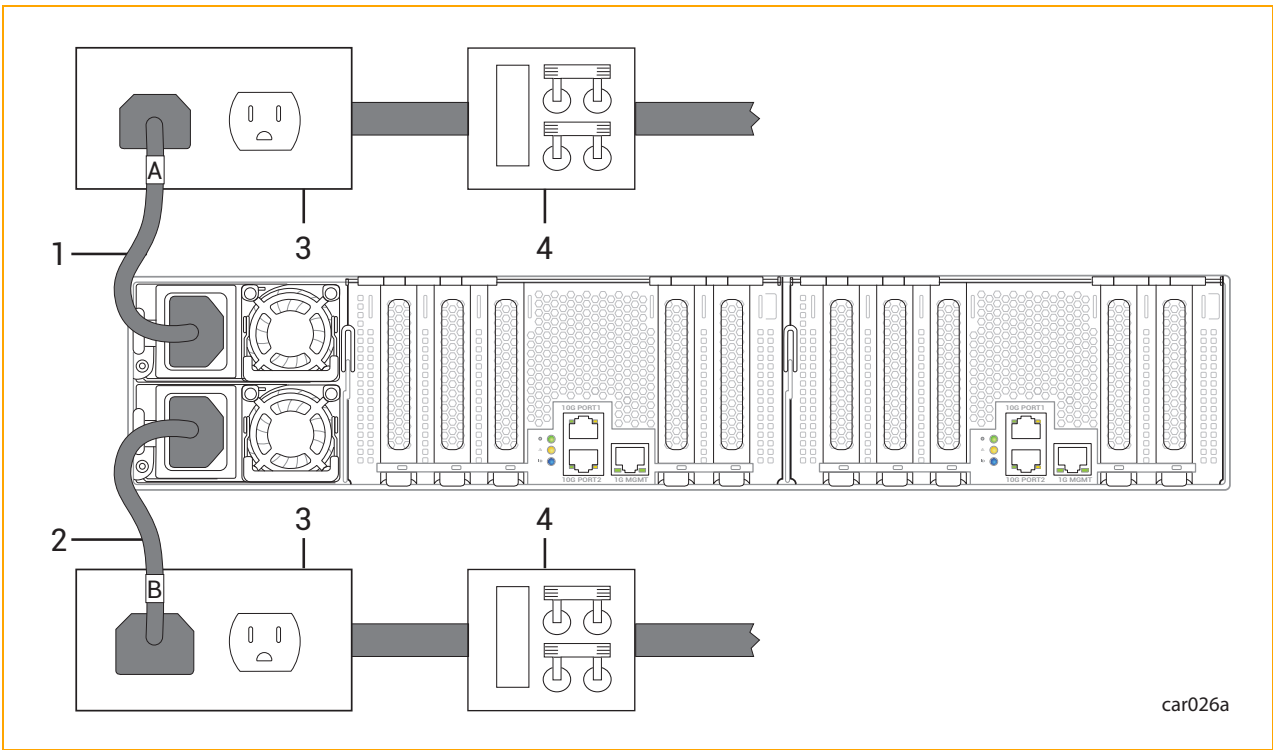
1	A-side power cord	3	AC power outlets
2	B-side power cord	4	AC power (mains) distribution circuit breakers (maximum of 20A)

Related Topics

Power Cords

Stratus supplies tested and approved AC power cords for the ztC Endurance system and for Stratus-supplied peripheral components.

> ⚠ **Caution**: Place all power cords out of the way of foot traffic.

> ⚠ **Précaution**: Éloigner tous les cordons d'alimentation du passage.

Power cords described as "Domestic" are made with American Wire Gauge (AWG) cordage and are intended for use in North America. Refer to UL 62 and CSA C22.2, No. 49. All other power cords are made with international cordage and are intended for use in the specified countries, and in other countries that use that standard.

Table 2-1 lists the available power cords that connect ztC Endurance systems directly to an AC power source (not to a PDU or to a UPS).

Table 2-1: Power Cords: ztC Endurance Systems to AC Power

Marketing ID	Plug Type	Rating	Length	Locale
ZEN-PWR-211	IEC 60320 C19 to CEE 7 VII	16A/250V	8.2 ft (2.5m)	Europe, Korea
ZEN-PWR-212	IEC 60320 C19 to IEC/EN 60309	16A/250V	8.2 ft (2.5m)	United Kingdom, Singapore
ZEN-PWR-213	IEC 60320 C19 to SAA/3/15	15A/250V	8.2 ft (2.5m)	Australia
ZEN-PWR-214	IEC 60320 C19 to IND/3/16 IS694	16A/250V	8.2 ft (2.5m)	India
ZEN-PWR-215	IEC 60320 C19 to JIS C8303	15A/250V	8.2 ft	Japan

	6/15AJ		(2.5m)	
ZEN-PWR-216	IEC 60320 C19 to GB2099/GB1002	16A/250V	8.2 ft (2.5m)	China
ZEN-PWR-217	IEC 60320 C19 to NBR14136	16A/250V	8.2 ft (2.5m)	Brazil
ZEN-PWR-218	IEC 60320 C19 to CNS 690	15A/250V	8.2 ft (2.5m)	Taiwan
ZEN-PWR-221	IEC 60320 C19 to NEMA 6-20	20A/250V	8.2 ft (2.5m)	United States
ZEN-PWR-225	IEC 60320 C19 to ZA/3 SABS 164-1	16A/250V	8.2 ft (2.5m)	South Africa
ZEN-PWR-226	IEC 60320 C19 to NEMA L6-20	20A/250V	8.2 ft (2.5m)	Japan
ZEN-PWR-220 (low voltage)	IEC 60320 C19 to 498GJTW	15A/125V	8.2 ft (2.5m)	Japan and Taiwan
ZEN-PWR-222 (low voltage)	IEC 60320 C19 to NEMA 5-20	20A/125V	8.2 ft (2.5m)	United States

Note: When ordering individual power cords as spares or upgrades, the product's Marketing ID should be followed by -P.

Related Topics

- "PSUs" on page 34
- "Removing and Replacing a PSU and Power Cords" on page 125
- Connecting the System to Electrical Power
- "Electrical Power Connections" on page 42

Grounding Requirements

Warning: Incorrect grounding can cause severe personal injury and extensive equipment damage.

> ⬣ **Avertissement**: Une mise à la terre incorrecte peut provoquer des blessures graves et endommager sérieusement l'équipement.

The ztC Endurance system obtains its protective earth (PE) ground through the power cord. See "Grounding Considerations" on page 131 for a more complete discussion of ztC Endurance system grounding.

Space Planning

For information about planning sufficient space for a ztC Endurance system, see:

- "Room Requirements" on page 45

- "Planning for Cables" on page 46

Room Requirements

To ensure that the installation site provides a properly equipped, cooled, and sized environment, make sure that the site:

- Provides clearances for air circulation, opening cabinet doors, removing cabinet panels, and servicing the system from the front and rear.

  Locate the front and rear of the system at least 2.5 feet (ft) (0.76 meters (m)) away from walls and other obstructions.

- Maintains reasonable temperature and humidity levels, and has a thermometer and humidistat to monitor room temperature and humidity.

  See "AC Power and HVAC Service Requirements" on page 41 for detailed information about HVAC planning and "System Specifications" on page 289 for detailed information about temperature and humidity requirements.

- Is as free as possible of airborne contamination (particulate and gaseous).

  > **Notice**: Due to the mission-critical nature of Stratus servers, data centers housing the equipment should follow the guidelines of airborne contamination (particulate and gaseous) as outlined in the ASHRAE (TC) 9.9 documentation. Particulate matter to adhere to ISO 14644-1 Class 8. Gaseous contamination, such as sulfur or chlorine-bearing gases to adhere to ANSI/ISA-71.04-1985 Severity level G1.

  Dust buildup in the system can impede air circulation and heat transfer, causing components to become less reliable as the ambient temperature rises.

  Fans clogged by dust fail to expel hot air, causing circuit boards to overheat and fail.

Dust on circuit boards raises the temperature, thus reducing the component's mean time between failure (MTBF).

Dust circulating in the room increases the risk of fire within the room by providing potential combustible material within the environment.

Dust contamination on tape devices causes mis-reads and -writes, leading to failure of attempts to back up and restore data.

- Provides two electrically separate grounded AC wall outlets, within reach of the power cords from the system or PDUs, and, if used, UPS units. If the power cord must be hardwired, consult an electrician to plan the length of the power cord.

- Provides cutouts in the floor for routing cables, if the site has an elevated floor.

- Contains space for future expansion.

> **Caution**: Do not place the system in an area of high electrostatic discharge. Static electricity may damage components. Do not locate components near transformers or other electromagnetic devices.

> **Précaution**: Ne pas placer le système dans une zone de forte décharge électrostatique. L'électricité statique peut endommager les composants. Ne pas placer de composants à proximité de transformateurs ou d'autres appareils électromagnétiques.

See for the dimensions of system components.

Planning for Cables

To accommodate cables from the system, make sure to provide:

- Ethernet jacks, switches, or hubs, as needed

- Two electrically separate grounded AC wall outlets, or a UPS and a wall outlet, within reach of the power cords from the system or PDUs, and additional outlets for any components that do not connect to a PDU

- AC wall outlets within reach of the power cords from optional components such as monitors, management PCs, or other users-supplied components; or use PDUs

Make sure that cables you plan to connect to the system are long enough to reach between the system and external components or connections. For information about specific cables and power cords, see the following topics:

-

-

-

Communications Line Planning

Make sure that you plan the locations of the ztC Endurance system and its external components so that all communications and data cables will reach their connection points.

> **Notes**:
>
> 1. If you have a unique network requirement, call Penguin Solutions Customer Service at the phone number listed on the the **Penguin Solutions Customer Service** page at https://www.penguinsolutions.com/en-us/support.
>
> 2. See "System Specifications" on page 289 for the PCIe adapters that Stratus supplies for the ztC Endurance system.

Related Topics

-

Cabinet Requirements

In general, consider the following factors during your site planning for systems in cabinets:

- Elevated Operating Ambient Temperature–If installed in a closed or multi-unit rack assembly, the operating ambient temperature of the rack environment may be greater than room ambient. Therefore, consideration should be given to installing the equipment in an environment compatible with the maximum ambient temperature (Tma) specified by the manufacturer.

- Reduced Air Flow–Installation of the equipment in a rack should be such that the amount of air flow required for safe operation of the equipment is not compromised.

- Mechanical Loading–Mounting of the equipment in the rack should be such that a hazardous condition is not achieved due to uneven mechanical loading.

- Circuit Overloading–Consideration should be given to the connection of the equipment to the supply circuit and the effect that overloading of the circuits might have on overcurrent protection and supply wiring. Appropriate consideration of equipment nameplate ratings should be used when addressing this concern.

- Reliable Earthing–Reliable earthing of rack-mounted equipment should be maintained.

Particular attention should be given to supply connections other than direct connections to the branch circuit (for example, use of power strips).

Make sure that the cabinet meets the following requirements:

- The cabinet can support the weight of all of the components you plan to install in it. "System Specifications" on page 289 provides the weights of ztC Endurance system components. To determine the weight of any peripheral components, see the documentation included with those components.

- The cabinet contains two front and two rear vertical EIA rails, one in each corner of the cabinet, that have the universal hole pattern as defined in the EIA 310-D Section 1 specification.

- The front vertical rails extend at least 0.5 in. (1.27 centimeters (cm)) beyond the inside edge of the accessory leg, if present, to allow the mounting rails to be fitted.

- The distance between the front vertical rails and the inside of the front door is at least 3.0 in. (7.62 cm).

- The distance between the rear of the system chassis and the inside of the rear door is at least 6.0 in. (15.24 cm).

- The distance between the front and rear vertical rails is between 24 in. and 31.5 in. (60.96 cm and 80 cm, respectively).

- The vertical mounting rails have a 0.275 in. (7.0 mm) round or 0.375 in. (9.53 mm) square pattern, with a thickness between 0.063 in. - 0.105 in.

- The vertical mounting rails are plated, or some other method is used to ensure continuity for grounding between installed equipment.

- Cable-management brackets are provided to support and constrain data and power cords so that the cables do not interfere with air flow out of the rear of the modules, and so that the connectors do not disconnect or break.

- To prevent stray voltages, all components are grounded together through the vertical mounting rails to the cabinet frame, and then to local building ground. To ensure signal quality, use a grounding cable for local building ground.

- There is a plan for maintaining cables and wires to the cabinet by either running them under the floor or placing them overhead in an overhead cable tray.

- Air flows through the cabinet from front to back.

- Air does not recirculate within the cabinet.

- Filler panels cover any unused rack space to prevent air recirculation.

- Blockers are installed between the vertical mounting rails and the side panels at the rear of the cabinet.

- Vents are evenly distributed on the front and rear doors and make up at least 69% of the surface area.

> **Caution**: The cabinet must have vented front and rear doors to prevent the system from overheating while in operation.

> **Précaution**: L'armoire doit avoir des portes avant et arrière ventilées pour empêcher le système de surchauffer pendant le fonctionnement.

- The final installation conforms to all emission, immunity, safety, and other applicable regulations.

> **Notice**: A ztC Endurance system allows you to install USB cables and USB flash drives in the front of the system, as well as various data and power cables in the rear of the system. The recommended clearance dimensions of 76 mm from front rails to the door, and 152 mm between the rear of the system chassis and the inside of the rear door, ensure that there are no clearance issues with the previously mentioned accessories or device. Although not recommended, a 1000 mm rack may be possible if the cabinet rails are skewed to the front, thus preventing use of accessories in the front of the system, while allowing more space for cabling in the rear.

Examples of recommended cabinets for system installation follow:

- With a vertical PDU installed, xU high x 31.5 in. (800 mm) wide x 47.2 in. (1200 mm) deep. For example: APC™ NetShelter™ SV AR2580-42U

- Without a PDU installed, xU high x 23.6 in. (600 mm) wide x 42.1 in. (1070 mm) deep. For example: APC NetShelter SX AR3104-24U

Site Planning for Peripheral Components

Remember to plan for any additional peripheral components, which can include a keyboard, mouse, and monitor, as well as other devices.

For your monitor, ensure that:

- The monitor accepts universal 100-240 volts AC (VAC), 50/60 Hertz (Hz) power.

- The VGA cable has a 15-pin D-sub connector if you are connecting it directly to a compute module.

- The power cord for the monitor is long enough to reach the power source.

- The plug type on the power cord is compatible with the external power source at the site.

ztC Endurance systems support the ZEN-PERIPH-KBMS keyboard and mouse. See "System Specifications" on page 289 for the specifications for this keyboard and mouse.

A 6-ft (1.83-m) USB cable is attached to the keyboard. Another 6-ft (1.83-m) USB cable is attached to the mouse.

3

Chapter 3: Installing and Connecting the ztC Endurance System Hardware

When you receive your ztC Endurance system, unpack it as shown on the flap of your system shipping box. You can then review ztC Endurance 3xx0, 5xx0, 7xx0, and 9xx0 Systems: Deploying a System (R002C), which is located in the shipping box, for a brief overview of your system.

> **Notice**: Make sure you read and understand the important safety information in "Safety Precautions" on page 51 before you begin installing and connecting the system.

Next Steps

Begin installing the ztC Endurance system as described in "Installing the System in a Cabinet" on page 62.

Safety Precautions

Before installing the system, take the following important precautions:

- Observe all applicable industry safety standards. See "Safety Notices" on page 52.

- Provide the necessary space and light to safely perform the installation.

- Do not wear conducting objects, such as rings, bracelets, and keys.

See "Warnings, Cautions, Notices, and Notes" on page 52 for information about the types of cautionary notes used in this document.

Warnings, Cautions, Notices, and Notes

Warnings, cautions, notices, and notes provide special information and have the following meanings:

> **Warning**: A warning indicates a hazardous situation that, if not avoided, could result in death or serious injury.

> **Avertissement**: Un avertissement indique une situation dangereuse qui, si pas évitée, pourrait entraîner la mort ou des blessures graves.

> **Caution**: A caution indicates a hazardous situation that, if not avoided, could result in minor or moderate injury.

> **Précaution**: Une mise en garde indique une situation dangereuse qui, si pas évitée, pourrait entraîner des blessures mineures ou modérées.

> **Notice**: A notice indicates information that, if not acted on, could result in damage to a system, hardware device, program, or data, but does not present a health or safety hazard.

> **Note**: A note provides important information about the operation of a ztC Endurance system or related equipment or software.

Safety Notices

In addition to the safety notices described here, also see:

- "Environmental Warnings" on page 57

- "Safety Precautions When Setting Up and Shipping Equipment" on page 58

- "Power Cord Warnings and Cautions" on page 58

- "Electrostatic Discharge (ESD) Warnings" on page 61

- "Rack System Installation" on page 62

Caution: Risk of explosion if battery is replaced by an incorrect type. Dispose of used batteries according to the instructions provided with the battery.

Précaution: Il y'a risque d'explosion si la batterie est remplacée per une batterie de type incorrect. Mettre au rebut les batteries usagées conformément aux instructions.

Warning: The system uses two power cords to provide redundant sources of power. To fully remove power from a system, disconnect both power cords.

Avertissement: Le système utilise deux cordons d'alimentation pour fournir des sources d'alimentation redondantes. Pour que le système ne soit plus alimenté, déconnecter les deux cordons d'alimentation.

Caution: To prevent a cabinet from tipping over and injuring a person or damaging the system, start installing systems from the bottom of the cabinet upward.

Précaution: Pour empêcher une armoire de basculer et blesser une personne ou endommager le système, installer le système sur la partie inférieure de l'armoire vers le haut.

Warning: To avoid fire, electric shock, and equipment breakdown, prevent water or foreign objects from getting into the equipment. Do not let water or foreign objects, such as pins or paper clips, enter the equipment.

Avertissement: Afin d'empêcher un incendie, un choc électrique et la panne de l'équipement, empêcher l'eau ou les objets étrangers de pénétrer dans l'équipement. Ne pas laisser l'eau ou les objets étrangers, comme des épingles ou des trombones, entrer dans l'équipement.

Warning: To prevent fire or current leakage, do not plug the power cord into a nonconforming outlet. Use a power outlet with appropriate voltage and power type, as specified in this guide.

Avertissement: Pour empêcher un incendie ou une fuite de courant, ne pas brancher le cordon d'alimentation dans une prise non conforme. Utiliser une prise de courant avec la tension et le type d'alimentation appropriés, comme indiqué dans ce guide.

Caution: Do not install the equipment where you may need an extension cord. Use of an extension cord that does not meet the power specifications introduces a risk of overheating that could lead to a fire.

Précaution: Ne pas installer l'équipement là où il est nécessaire d'utiliser une rallonge. L'utilisation d'une rallonge non conforme aux spécifications en matière d'alimentation induit un risque de surchauffe, ce qui pourrait causer un incendie.

Warning: Disconnect the power cords from the server or power source before you install or relocate the equipment. All voltage is removed only when the power cords are disconnected.

Avertissement: Débrancher les cordons d'alimentation du serveur ou de la source d'alimentation avant d'installer ou de déplacer l'équipement. La tension est éliminée uniquement lorsque les cordons d'alimentation sont débranchés.

Warning: To reduce the risk that electrical shock could injure a person or damage the system, exercise caution when working in the unit even when only one power cord is connected.

Avertissement: Pour réduire le risque qu'un choc électrique puisse blesser une personne ou endommager le système, soyez prudent lorsque vous travaillez dans l'appareil, même lorsqu'un seul cordon d'alimentation est connecté.

Warning: Do not disassemble, repair, or alter the server, except as described in this documentation. There are no user-serviceable parts within the chassis. The modules and PSU are all tool-less serviceable. Do not use any tools to access or remove any panels or covers. There is a risk of an electric shock or fire as well as equipment malfunction if you do not observe the instructions in this documentation.

Avertissement: Ne pas démonter, réparer, ou altérer le serveur, excepté tel que décrit dans cette documentation. Il n'y a pas de pièces pouvant être réparées par l'utilisateur dans le châssis. Les modules et le bloc d'alimentation peuvent être réparés entièrement sans outil. Ne pas utiliser d'outils pour retirer ou accéder aux panneaux ou aux couvercles. Il y a un risque de choc électrique ou d'incendie, ainsi que de mauvais fonctionnements de l'équipement, si vous ne respectez pas les instructions dans cette documentation.

Caution: Do not place any object on top of the server. The object may fall off and cause injuries, damage to hardware, or a fire.

Précaution: Ne pas placer d'objet en haut du serveur. L'objet peut chuter et causer des blessures, endommager le matériel, ou causer un incendie.

Caution: To reduce the risk of personal injury from hot surfaces, allow the drives and the internal system components to cool before touching them.

Précaution: Pour diminuer le risque de blessure au contact de surfaces chaudes, laissez refroidir les lecteurs et les composants internes du système avant de les toucher.

Warnings:

To reduce the risk of personal injury or damage to the equipment:

1. Observe local occupational health and safety requirements and guidelines for manual material handling.

2. Get help to lift and stabilize the product during installation or removal, especially when the product is not fastened to the rails. This server weighs between 74-102 lb (33.6-46.3 kg).

Avertissements:

Pour diminuer le risque de blessure ou de dommage à l'équipement:

1. Respectez les exigences et les directives en matière de santé et de sécurité au travail pour la manutention manuelle du matériel.

2. Faites-vous aider pour lever et stabiliser le produit pendant l'installation ou le retrait, en particulier quand le produit n'est pas fixé aux rails. Le serveur pèse entre 33,6 et 46,3 kg (74 à 102 lb). Ce poids indique qu'il faut au minimum trois personnes pour lever le système pendant le déballage, le déplacement et l'installation du système dans un bâti.

Caution: Observe caution when inserting compute, storage, or I/O modules. These are moving parts. Keep fingers clear to avoid any possible pinching when pushing the modules into a latched position.

Précaution: Soyez prudent lorsque vous insérez une unité centrale de traitement, de la mémoire ou des modules E/S. Ces pièces sont mobiles. Garder les doigts à l'écart pour éviter un pincement au moment de verrouiller les modules en place.

Notice: Do not install or store the equipment in an unsuitable place. Install or store the equipment in a place that meets the requirements specified in this guide.

Environmental Warnings

Cautions:

1. Do not use or store this product in a corrosive environment.

2. Avoid using or storing this product in an environment that may contain corrosive gases. Such gases include, but are not limited to, sulfur dioxide, hydrogen sulfide, nitrogen dioxide, chlorine, ammonia, or ozone.

3. Avoid installing this product in a dusty environment or one that may contain corrosive materials such as sodium chloride or sulfur.

4. Avoid installing this product in an environment having excessive metal flakes or conductive particles in the air. Such environments may cause corrosion or short circuits within this product, resulting in not only damage to this product, but also fire.

5. If there are any concerns regarding the environment at the planned site of installation or storage, contact your authorized Penguin Solutions service representative.

Précautions:

1. Ne pas utiliser ou entreposer ce produit dans un milieu corrosif.

2. Éviter d'utiliser ou d'entreposer ce produit dans un environnement qui peut contenir des gaz corrosifs. Ces gaz incluent, mais sans s'y limiter, le dioxyde de soufre, le sulfure d'hydrogène, le dioxyde d'azote, le chlore, l'ammoniaque ou l'ozone.

3. Éviter d'installer ce produit dans un environnement poussiéreux ou dans un environnement qui contient des matières corrosives comme le chlorure de sodium ou le soufre.

4. Éviter d'installer ce produit dans un environnement qui contient des écailles de métal ou des particules conductrices dans l'air. Ce type d'environnement peut causer de la corrosion ou des courts-circuits dans le produit, ce qui endommagerait le produit, mais causerait également un incendie.

5. Si vous avez des inquiétudes concernant l'environnement du site où il est prévu d'installer ou entreposer le produit, contactez votre représentant de service Penguin Solutions autorisé.

Safety Precautions When Setting Up and Shipping Equipment

For your protection, observe the following safety precautions when setting up and shipping your equipment:

- Follow all cautions and instructions marked on the equipment.

- Ensure that the voltage and frequency of your power source match the voltage and frequency written on the equipment's electrical rating label.

- Never push objects of any kind through openings in the equipment. Dangerous voltages, energy, or moving parts may be present. Conductive external objects could produce a short circuit that could cause fire, electric shock, or damage to your equipment.

- Do not block or cover the openings of your product. Never place a product near a radiator or heat register. Failure to follow these guidelines can cause overheating and affect the reliability of your product.

- Do not drop the product or subject it to physical shock.

- Keep the product away from water or any liquid.

- To ship the product, pack it inside the qualified package and ship with a pallet.

Power Cord Warnings and Cautions

Installation of this equipment must comply with local and regional electrical regulations governing the installation of information technology equipment by licensed electricians. For electrical power ratings on options, refer to the power rating label or the user documentation supplied with that option.

Warnings:

1. When installing a system, always connect the power cord first, before adding communications cables.

2. The power cord contains the protective earth connection; it should be connected first and disconnected last, to maintain a grounded chassis.

3. Before attempting to remove a module from the system chassis, power off the module, disconnect communications cables, and then disconnect the power cord. Before attempting to remove a PSU, disconnect the power cord.

4. Never connect a power cord to a module or PSU when it is not located within the system chassis.

5. Do not connect communications cables to an I/O module that is not installed in a grounded system chassis. Disconnect all communications cables to the I/O module before removing the I/O module from the chassis.

Avertissements:

1. Lors de l'installation du système, toujours connecter le cordon d'alimentation en premier, avant d'ajouter les câbles de communication.

2. Le cordon d'alimentation inclut le raccordement à la terre, il doit être connecté en premier et déconnecté en dernier afin que le châssis reste relié à la terre.

3. Avant d'essayer de supprimer un module du châssis du système, s'assurer que le module est hors tension et les câbles de communication sont débranchés. Avant d'essayer de retirer un bloc d'alimentation, débrancher le cordon d'alimentation.

4. Ne jamais connecter un cordon d'alimentation à un module ou un bloc d'alimentation quand il n'est pas situé dans le châssis du système.

5. Ne pas connecter les câbles de communication à un module E/S qui n'est pas installé dans le châssis d'un système relié à la terre. Débrancher tous les câbles de communication du module E/S avant de retirer le module E/S du châssis.

Cautions:

1. This product is designed to work with multiple types of grounding systems, specifically TN, where neutral is grounded and for both IT and TN-C where grounded neutral is not used. For IT and TN-C, the branch circuit for L1 and L2 must be supplied by a double-pole circuit breaker. Contact your facilities manager or a qualified electrician if you are not sure what type of power is being supplied to the system.

2. Not all power cords have the same current ratings. Do not use the power cord provided with your equipment for any other products or use. Do not use household extension cords with your product.

3. The outlet for the DC power supply should not be far away from the equipment when using.

Précautions:

1. Ce produit est conçu pour fonctionner avec plusieurs types de systèmes de mise à la terre, en particulier TN, où le neutre est mis à la terre, et avec les systèmes IT et TN-C où le neutre relié à la terre n'est pas utilisé. Pour les systèmes IT et TN-C, le circuit de dérivation pour L1 et L2 doit être fourni par un disjoncteur à deux pôles. Contactez votre responsable de réseau ou un électricien qualifié si vous n'êtes pas sûr de savoir quel type d'alimentation est fournie au système.

2. Tous les cordons d'alimentation n'ont pas les mêmes courants nominaux. Utiliser uniquement le cordon d'alimentation fourni avec votre équipement. Ne pas utiliser des cordons d'alimentation domestiques avec votre produit. Ne pas utiliser le cordon d'alimentation fourni avec votre équipement avec d'autres produits ou pour une autre utilisation.

3. La prise d'alimentation C.A. ne doit pas être éloignée de l'équipement pendant son utilisation.

Electrostatic Discharge (ESD) Warnings

Cautions:

1. In order to maximize overall product quality and eliminate the possibility of damage from electrostatic discharge, observe ESD handling procedures wherever this symbol occurs during the system integration.

2. All employees must attend an ESD training class prior to performing any system or board handling. If you have not been properly trained, contact your immediate supervisor to schedule the training.

3. Use heel straps on both feet and/or wrist straps where applicable to ensure that you are properly grounded. Test the continuity of the straps at a certified ESD station before working with the products.

4. If you remove your heel straps or wrist straps for any reason, you must retest to ensure proper grounding.

5. Ensure the policies regarding ESD controls for all personnel who work in the manufacturing departments are enforced.

Précautions:

1. Afin de maximiser la qualité générale du produit et d'éliminer la possibilité de dommage dû à une décharge électrostatique, respectez les procédés de manutention ESD quand le symbole suivant apparaît pendant l'intégration du système.

2. Tous les employés doivent suivre un cours de formation ESD avant d'effectuer la manutention du système ou du tableau. Si vous n'avez pas été correctement formé, contactez votre superviseur immédiat pour planifier la formation.

3. Des bandes antistatiques pour talon à chaque pied et/ou des bracelets antistatiques, le cas échéant, doivent être utilisés pour garantir que vous êtes correctement mis à la terre pendant le test de continuité des bandes au poste ESD certifié.

4. Si vous retirez vos bandes antistatiques pour talon ou vos bracelets antistatiques pour une raison particulière, vous devez refaire le test pour garantir la mise à la terre correcte.

5. Assurez-vous que les politiques concernant les contrôles ESD sont appliquées pour tout le personnel qui travaille dans les départements de production.

Rack System Installation

The installation instructions include rack-mount instructions identical to or similar to the ones listed in "Cabinet Requirements" on page 47. Be sure to carefully read and observe the Cautions listed in that topic.

Related Topics

- "General Safety Precautions" on page 105

Installing the System in a Cabinet

The general steps for installing the ztC Endurance system in a cabinet are as follows:

1. Number the vertical rails

2. Prepare the support rails for installation

3. Attach the support rails and insert the system into the cabinet

However, before installing the system in a cabinet, make sure that:

- You read and understand all safety precautions in "Safety Precautions" on page 51.

- You read and understand all of the installation steps.

- Your cabinet meets the system requirements, and the site has been properly prepared (see "Cabinet Requirements" on page 47).

- You write down the asset ID, which is located on top of the system chassis (see Figure 3-1), and store it in a convenient location.

> **Caution**: Making note of the asset ID now is critical. After you install the system, the asset ID label may not be visible in the cabinet, and you need to provide the asset ID during the software installation.

Figure 3-1: Location of Asset ID

car038a

1	Label with asset ID
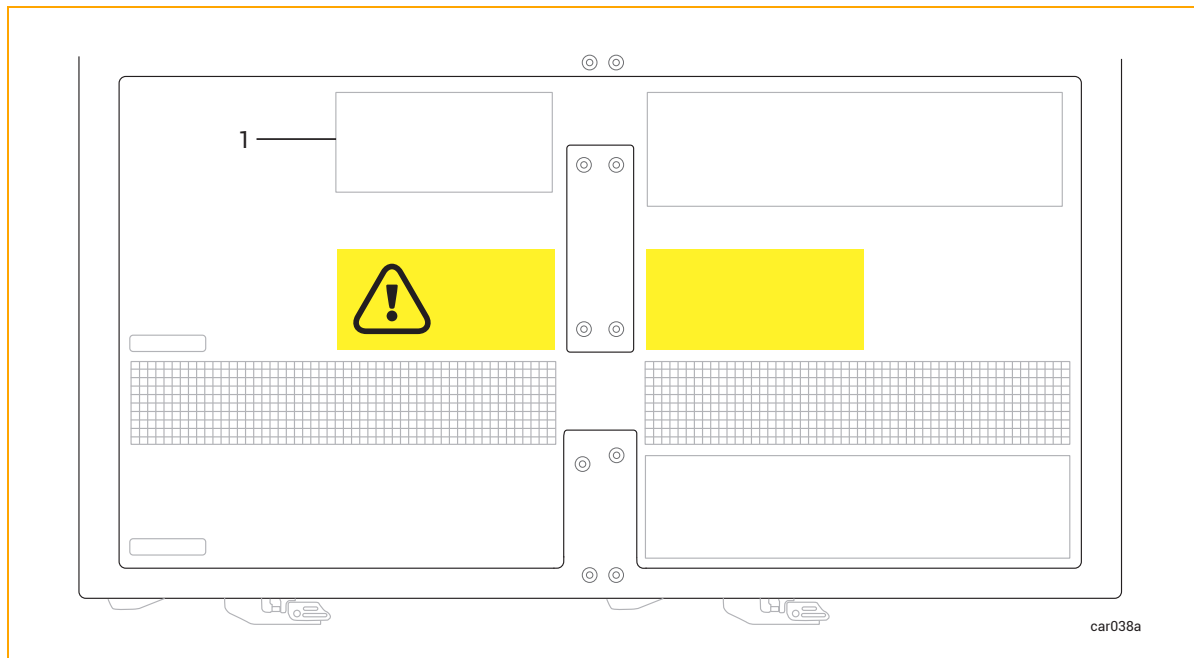
- You have the components to install:

  - Support rails (2). One is labeled **FRONT LEFT**, and the other is labeled **FRONT RIGHT**. The rails can be extended from 26 3/8 in. (670 mm) to 34 ¼ in. (870 mm).

  - Package labeled **R0897-E0112-01** that contains eight flat-head alignment screws (silver). Use these screws to align the mounting rails to the cabinet during installation.

  - Package labeled **R0897-E0113-01** that contains four Phillips/hex-head screws (black). Use these screws to secure the rail to the rear of the cabinet.

- You have the required tools: a flat-head screwdriver and a #2 Phillips-head screwdriver.

> **Caution**: To avoid personal injury or damage to the system, two persons are recommended to insert the system in a cabinet.

> **Précaution**: Pour éviter toute blessure corporelle ou tout dommage à l'étagère, deux personnes sont recommandées pour insérer l'étagère dans l'armoire et sur les supports de montage.

Next Steps

Start installing the system in a cabinet as described in "Numbering the Vertical Rails" on page 64.

Numbering the Vertical Rails

When you install a system in a cabinet, you first need to number the vertical rails, if they are not already numbered. To do so, use a marker or pencil to number the holes of the 2U in which to install each system.

The cabinet rails, system, and various components are measured in Electronic Industries Alliance (EIA) rack units (U). Every three holes on the EIA square-hole pattern on the vertical rail equals one U (1.75 in., 4.45 cm). The holes on EIA rails are in a repeating pattern of 1/2 in. (1.27 cm), 5/8 in. (1.59 cm), 5/8 in. (1.59 cm), as shown in Figure 3-2. The middle hole of each U has equal spacing between it and the holes adjacent to it. In some cabinets, numbers on the vertical rails identify the middle hole of each U.

Figure 3-2: EIA Vertical Rail Hole Pattern
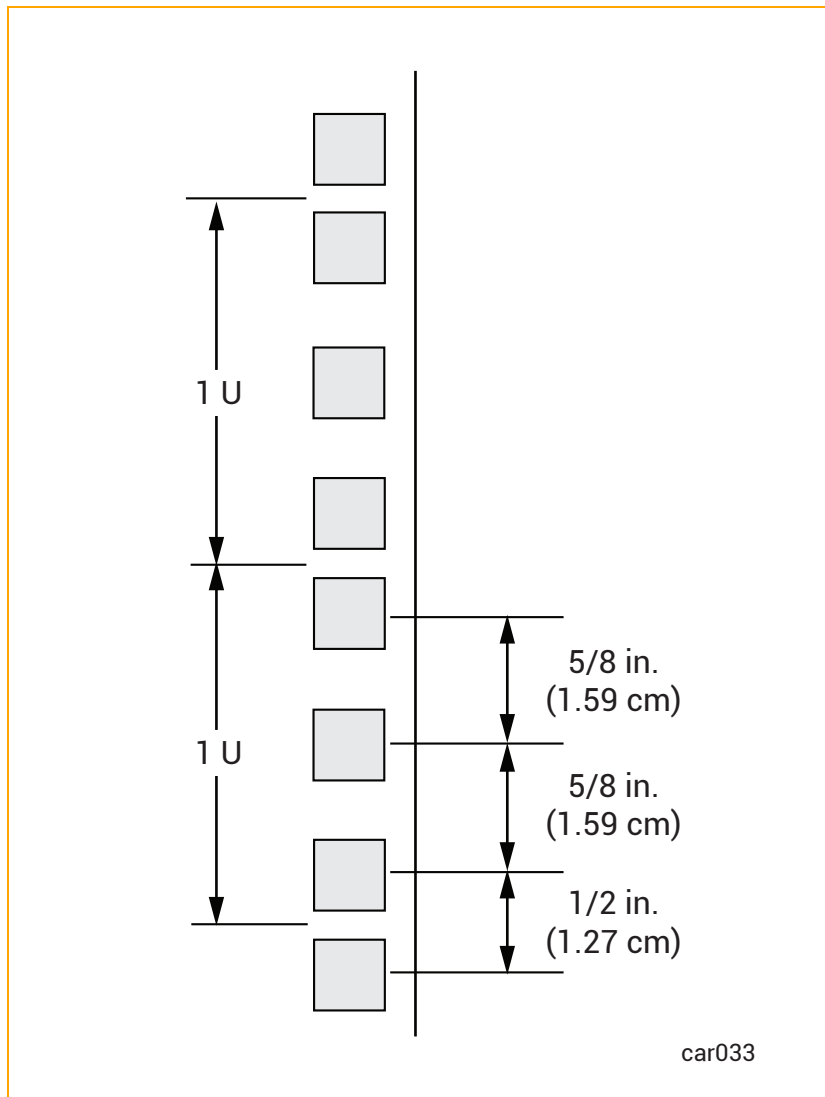
To number the vertical rails

1. On each of the front and rear vertical rails, number the 2U spaces where you will install the system.

   > **Notice**: To ensure that the components will be level when you install them, mark the same holes on the rear vertical rails that you marked on the front vertical rails.

2. If you are installing additional systems in the same cabinet, continue numbering another 2U for each system.

Next Steps

To continue with installing a system in a cabinet, see "Preparing the Support Rails for Installation" on page 66.

Preparing the Support Rails for Installation

When installing a system in a cabinet, you need to prepare the support rails for installation after you have numbered the vertical rails.

To prepare the support rails

1. Unpack the components and place them on a stable work surface. Attached to the rails are alignment/support screws for use in cabinet rails with round 13/64 in. (5 mm) holes.

   - If your cabinet rails have round 13/64 in. (5 mm) mounting holes, skip this section and continue to "Attaching the Support Rails and Inserting the System" on page 69.

   - If your cabinet rails have square 3/8 in. (9.525 mm) mounting holes, use the silver alignment screws (located in the package labeled **R0897-E0112-01**) in place of the screws that are already installed. Go to step 2.

2. Positioning a flat-head screwdriver in the middle of the front latch, pry open the latch. See Figure 3-3.

Figure 3-3: Positioning Screwdriver and Opening Latch

car010
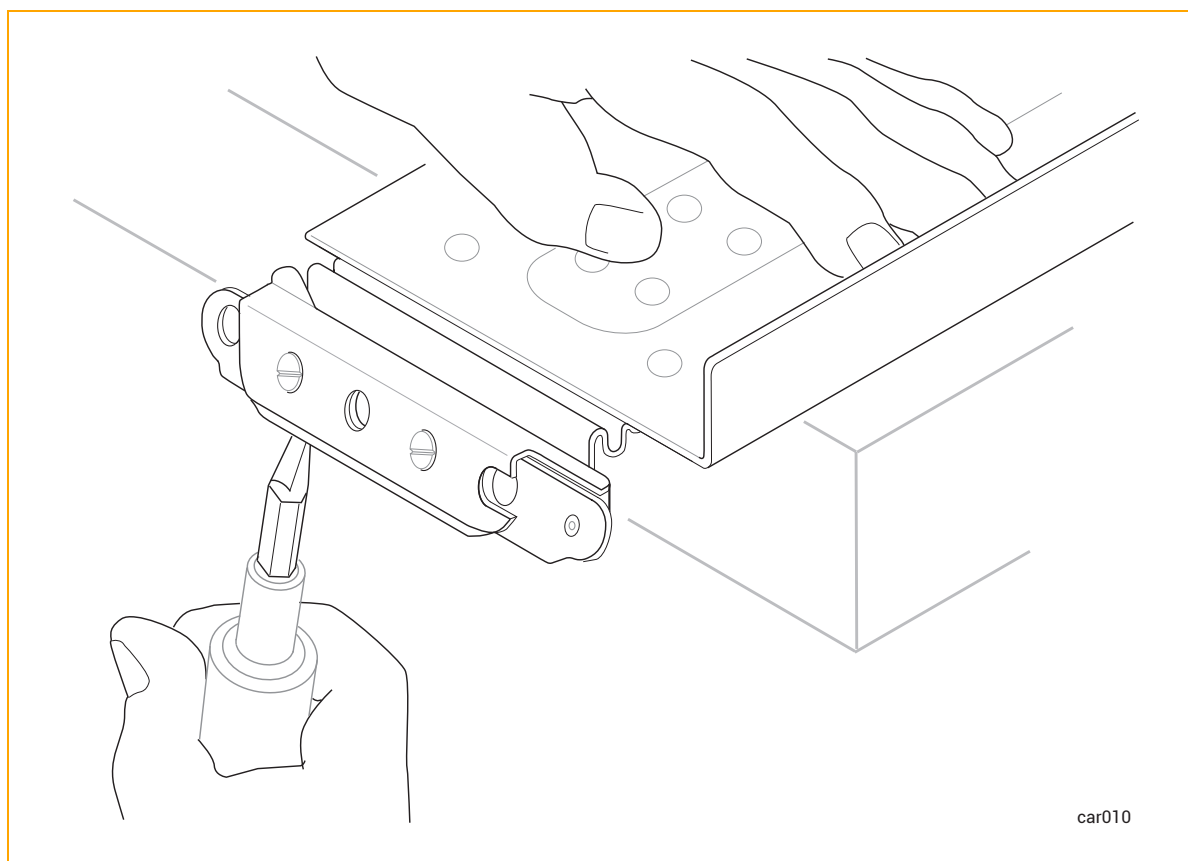
3. While the latch is in the open position, use a flat-head screwdriver to remove the two existing alignment screws. See Figure 3-4.

Figure 3-4: Removing Existing Alignment Screws

car011
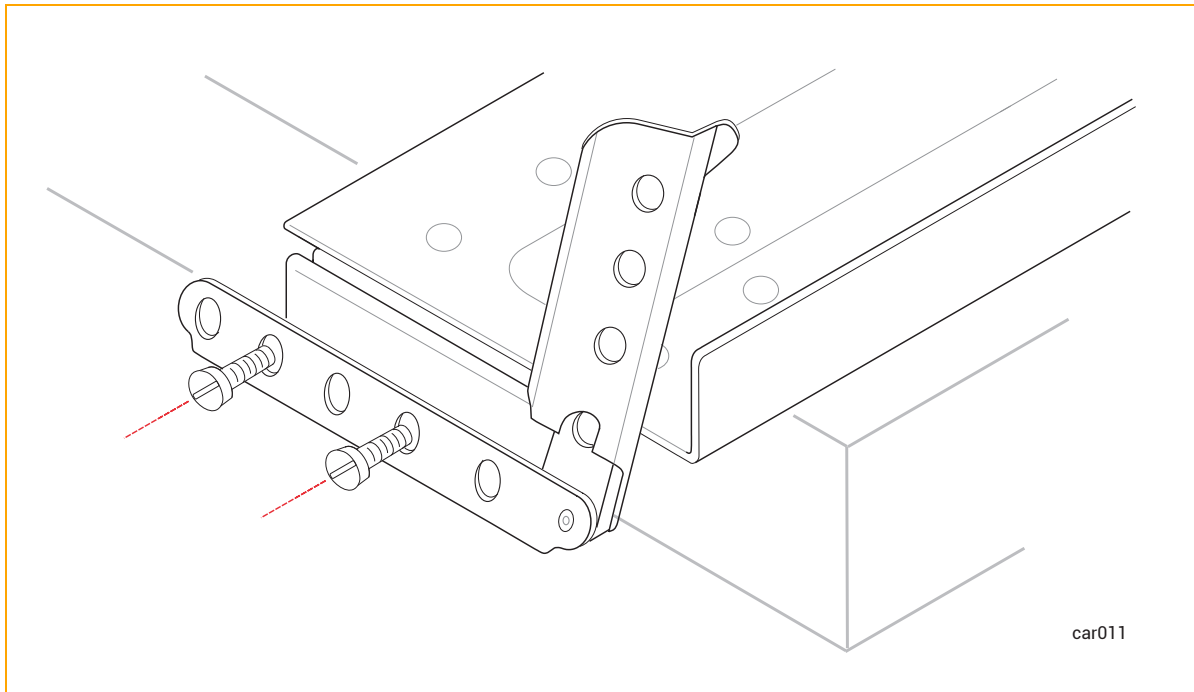
4. Use a flat-head screwdriver to attach two silver alignment screws in the holes that the previous alignment screws were in. Keep the latch open. See Figure 3-5.

Figure 3-5: Attaching Silver Alignment Screws

car012
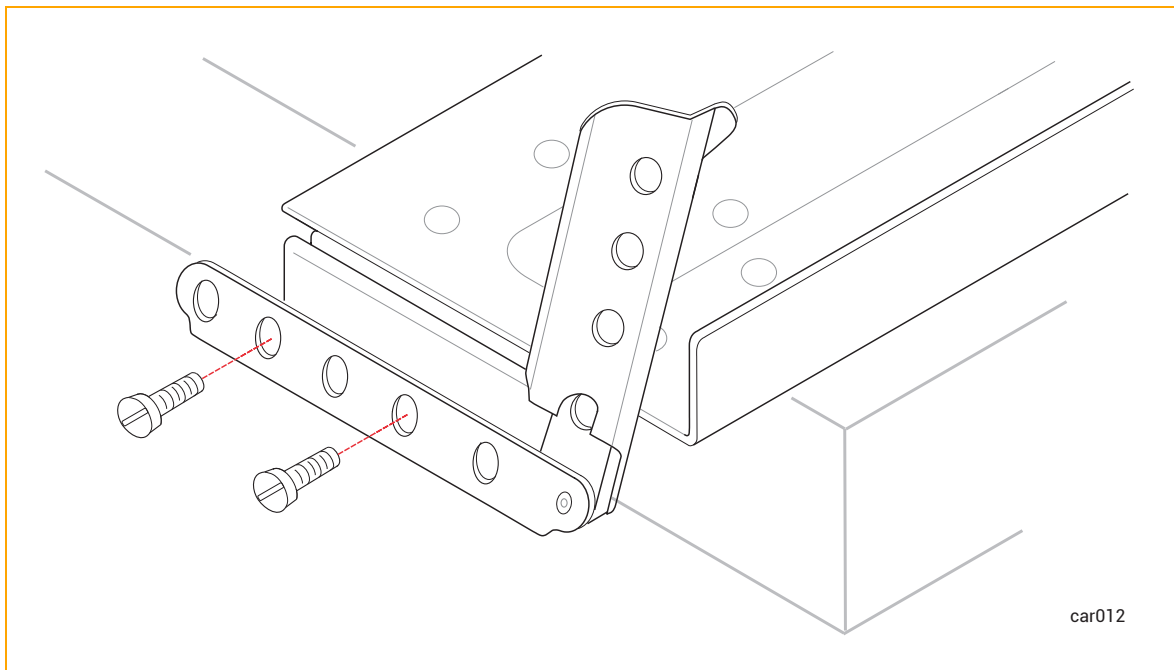
5. Repeat steps 3-4 for the rear latch.

6. Repeat steps 3-5 for the other support rail.

Next Steps

To continue installing a system in a cabinet, see "Attaching the Support Rails and Inserting the System" on page 69.

Attaching the Support Rails and Inserting the System

After you have prepared the support rails for installation, attach the support rails to the vertical rails in the cabinet, and then insert the system into the cabinet.

To attach the support rails

1. Hold the support rail labeled **FRONT LEFT** and face the front of the cabinet.

2. Align the silver alignment screws in the front of the left support rail with the 2U holes that you marked in "Numbering the Vertical Rails" on page 64 (that is, the third hole of the first U and the second hole of the second U). Close the latch. See Figure 3-6.

Figure 3-6: Securing Front of Support Rail

car013
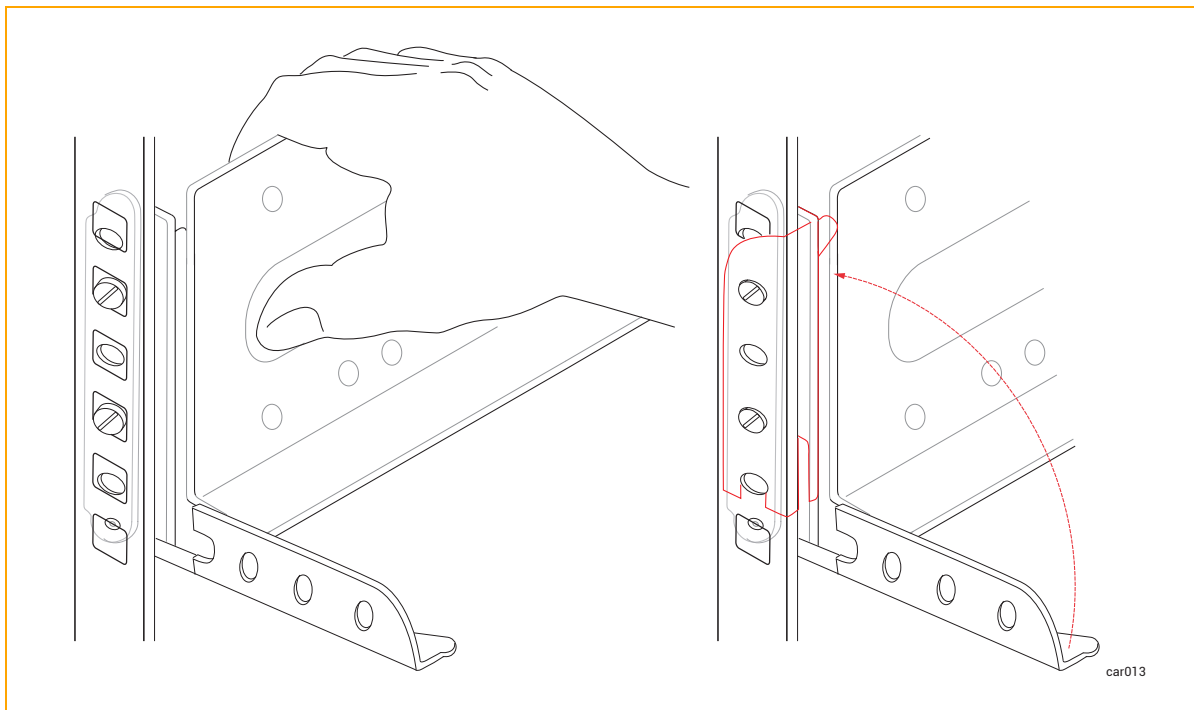
3. Moving to the rear of the cabinet, align the silver alignment screws in the rear of the left support rail with the 2U holes that you marked in "Numbering the Vertical Rails" on page 64. Adjust the

support rail as necessary to fit the depth of the cabinet and close the latch.

4. Secure the rear of the rail by using a Phillips-head screwdriver to attach two black Phillips-head/hex screws in the holes below the silver alignment screws. See Figure 3-7.

Figure 3-7: Securing Rear Support Rail

car014
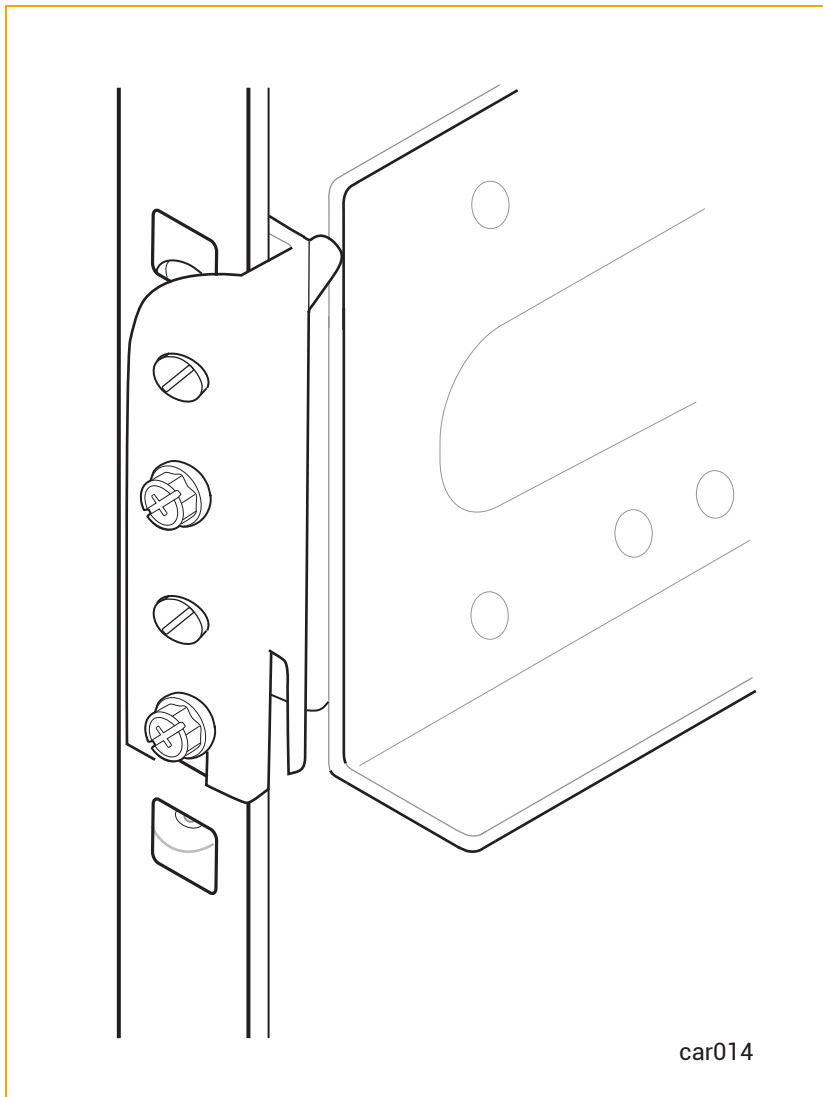
Perform steps 1-4 for the right support rail.

To insert the system in a cabinet

> **Notice**: Remove the compute modules before you install the system in a cabinet. Doing so reduces the overall weight of the system, making the system easier to lift and install.

> ⚠ **Caution**: To avoid personal injury or damage to the system, up to three persons may be needed to insert the system in a cabinet.

> ⚠ **Précaution**: Pour éviter toute blessure corporelle ou tout dommage à l'étagère, au moins trois personnes sont fortement recommandées pour insérer l'étagère dans l'armoire et sur les supports de montage.

1. With one person on each side of the system and (ideally) one person in the front of the system, align it in the front cabinet opening and place it on the mounting rails.

2. Slide the system into the cabinet until the left and right ears are flush with the cabinet rails. See Figure 3-8.

Figure 3-8: Inserting System into Cabinet

car015a
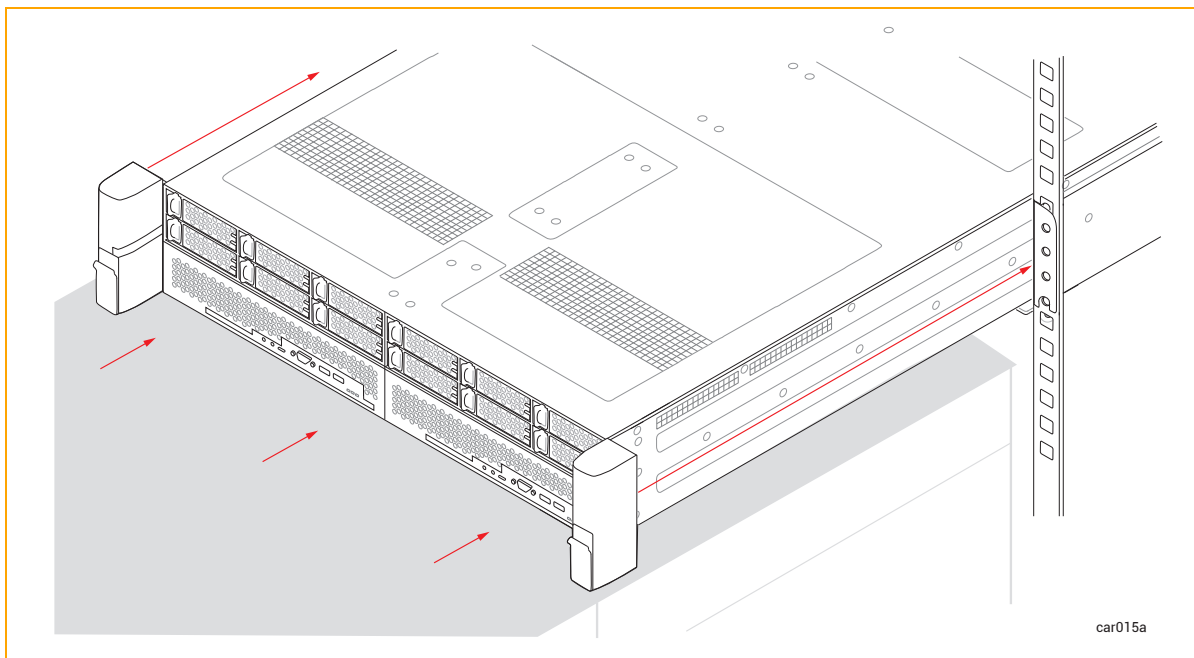
3. On the right ear, pull down the latch and use the Phillips-head screwdriver to secure the captive screw to the rail. Securing this captive screw fastens the system chassis and the front rail to the vertical rail of the cabinet. See Figure 3-9.

Figure 3-9: Attaching Captive Screws

car027
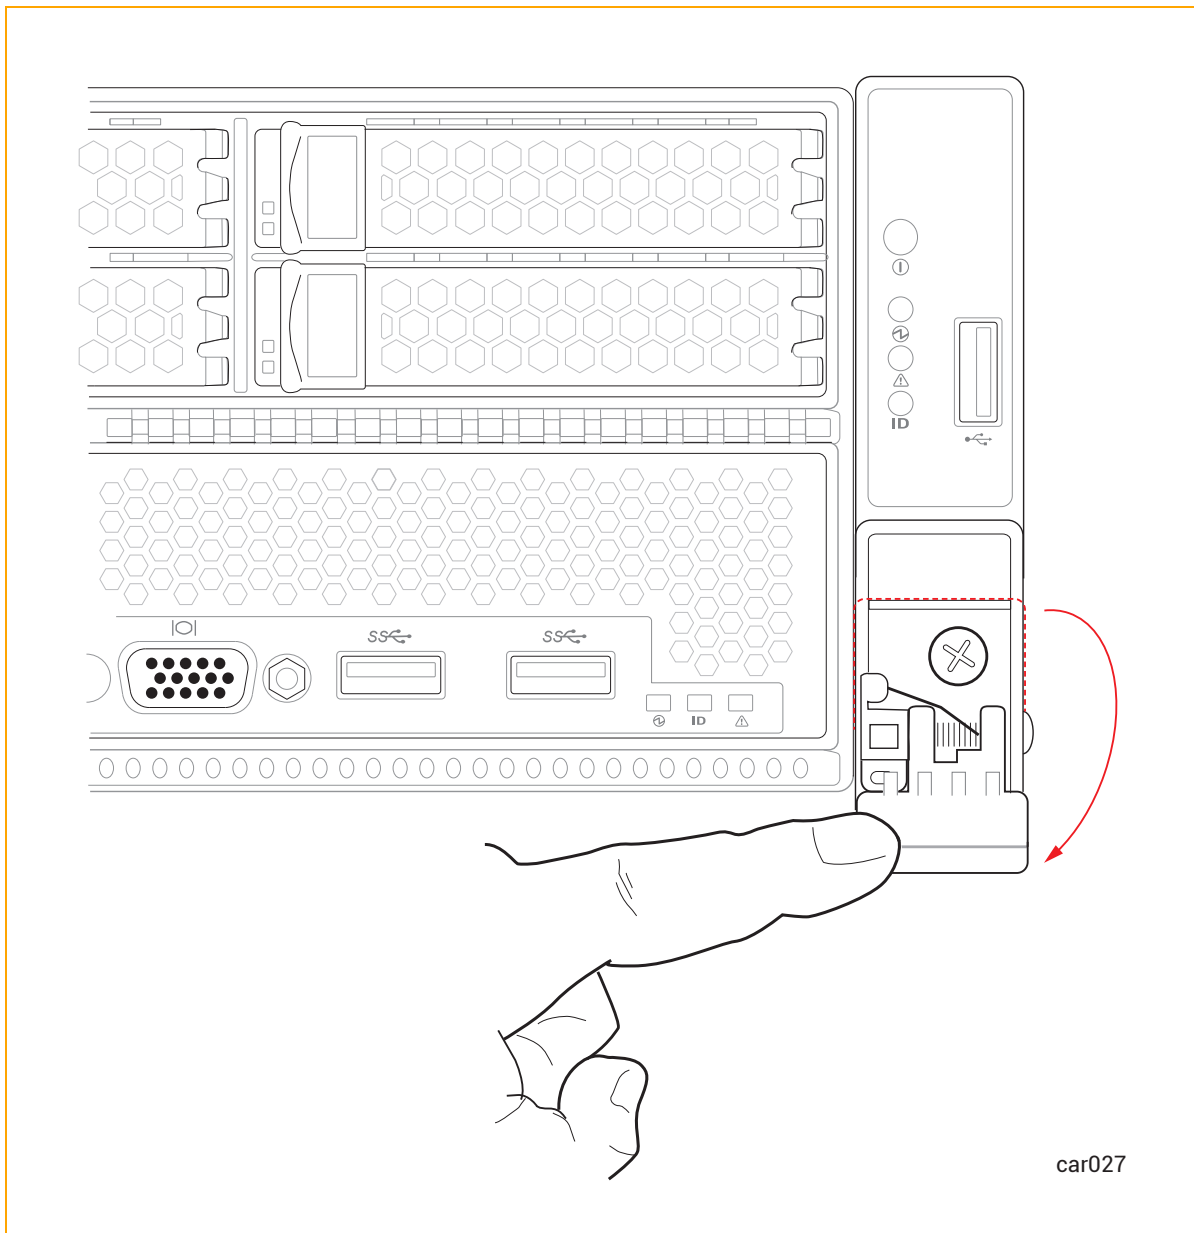
4. Repeat step 3 for the left ear.

Next Steps

Complete the hardware set-up, as described in "Completing the Hardware Set-up" on page 73.

Connecting the System Components

To connect your system components, see:

- "System Ports" on page 19 for illustrations of the ports on a ztC Endurance system

- "I. Connect Network Cables" on page 74 for information about connecting network cables

- Connecting Peripherals for information about connecting a keyboard, mouse, and monitor

- Connecting the System to Electrical Power for information about connecting the system to electrical power

No tools are required for these tasks.

Completing the Hardware Set-up

After you install the ztC Endurance system hardware, complete the system hardware set-up:

"I. Connect Network Cables" on page 74

"II. Connect a Mouse, Keyboard, and Monitor" on page 75

"III. Connect the System to Electrical Power" on page 77

"IV. Insert Blank Disks into Storage Modules" on page 83

No tools are required for these tasks.

> **Note**: When routing cables within the cabinet, keep all power cords as separate as possible from peripheral and communication cables. For example, route all power cords on one side of the cabinet and all peripheral and communications cables on the other side of the cabinet, using the cable-management rings on your cabinet to constrain the cords and cables.

I. Connect Network Cables

In both I/O modules, connect network cables from the bottom two **10G PORT2** ports to a production network and from the **1G MGMT** port to a management network. Figure 3-10 shows the location of the network ports on the rear of the system.

Figure 3-10: Location of Network Ports

1	10G ports (4)
2	1G MGMT ports (2)
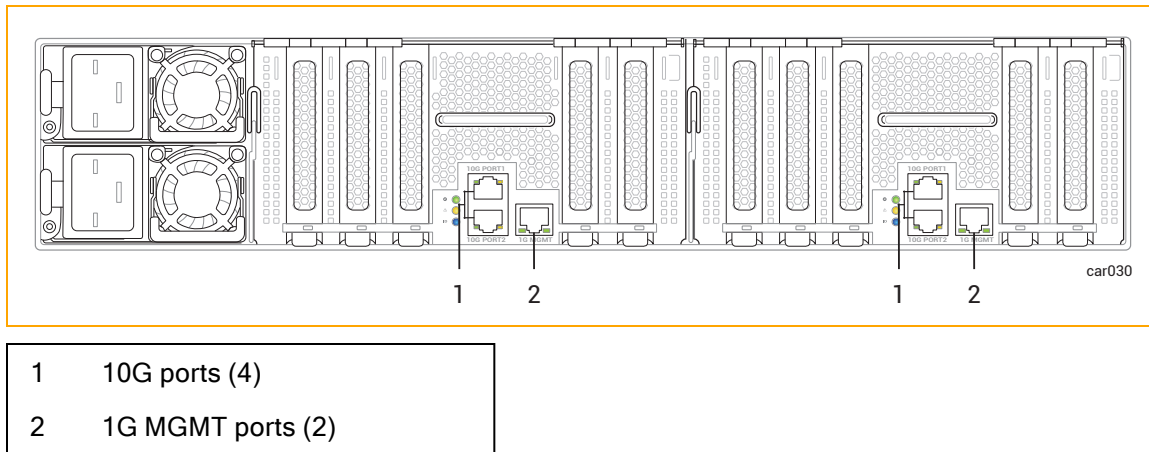
II. Connect a Mouse, Keyboard, and Monitor

To prepare for the software installation, connect two monitors (user-supplied) and two keyboards and mice (supplied by Stratus, if ordered) to ports on the system.

> **Notes**:
>
> 1. Because the installation requires you to respond to prompts on both compute modules, connecting two monitors and keyboards means that you can view the activities on both compute modules simultaneously. If you can connect only one monitor and keyboard, you need to open a separate KVM console for each compute module.
>
> 2. Other keyboards and mice have not been tested and may not work properly.

For information about monitor and keyboard requirements, see "Site Planning for Peripheral Components" on page 49.

To connect a mouse, keyboard, and VGA monitor

1. Connect the USB cables from the mouse and keyboard to USB ports on a compute module (see Figure 3-11).

> **Notice**: Do not connect the mouse or keyboard to the USB 2.0 port on the control-panel ear.

Figure 3-11: Location of VGA and USB Ports

car029

1	Debug ports (2)	3	USB 3.0 ports (4)
2	VGA (monitor port) (2)	4	USB 2.0 port (1)
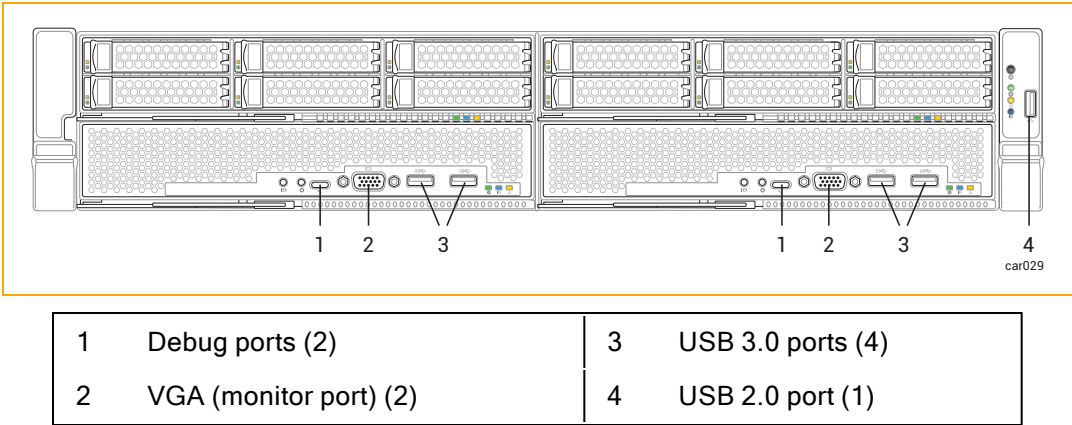
2. Connect the 15-pin D-sub connector on the VGA cable from the monitor to the system's VGA port on a compute module (see Figure 3-11).

3. Connect the VGA cable to the monitor.

III. Connect the System to Electrical Power

For both power supply units (PSUs), connect power cords from the power inlets to supported power sources, and then fasten the power-cord retention strap. This strap helps prevent unintentionally disconnecting the power cord from the system if the system is removed from the cabinet.

For information about electrical power planning, see "Electrical Power Planning" on page 40.

> **Warnings**:
>
> 1. A ztC Endurance 31x0 or 51x0 system can operate in either the low-voltage or high-voltage range (see "System Specifications" on page 289 for details). Make sure you are using power cords that have the correct voltage range and that are rated for your country. If you are changing to a different voltage range, follow the instructions described in To replace both power cords.
>
> 2. Both PSUs in a ztC Endurance system must operate in the same voltage range.

> **Avertissements**:
>
> 1. Le système ztC Endurance 31x0 ou 51x0 peut fonctionner en basse ou haute tension (voir les spécifications du système pour plus de détails). Assurez-vous que vous utilisez des cordons d'alimentation dont la plage de tension est correcte et qui sont adaptés à votre pays. Si vous passez à une plage de tension différente, suivez les instructions décrites dans la section Pour remplacer les deux cordons d'alimentation.
>
> 2. Les deux blocs d'alimentation d'un système ztC Endurance doivent fonctionner dans la même plage de tension.

To connect the system to AC power mains

1. From the front and back of the system, make sure that all modules and PSUs are fully inserted.

2. Locate the two power cords for the system.

3. Connect the female end of the A-side power cord to the top PSU, as shown in Figure 3-12. Connect the other end to the A-side power source.

4. Connect the female end of the B-side power cord to the bottom PSU, as shown in Figure 3-12. Connect the other end to a second, separate, B-side power source.

Figure 3-12: Connecting a System Directly to AC Power
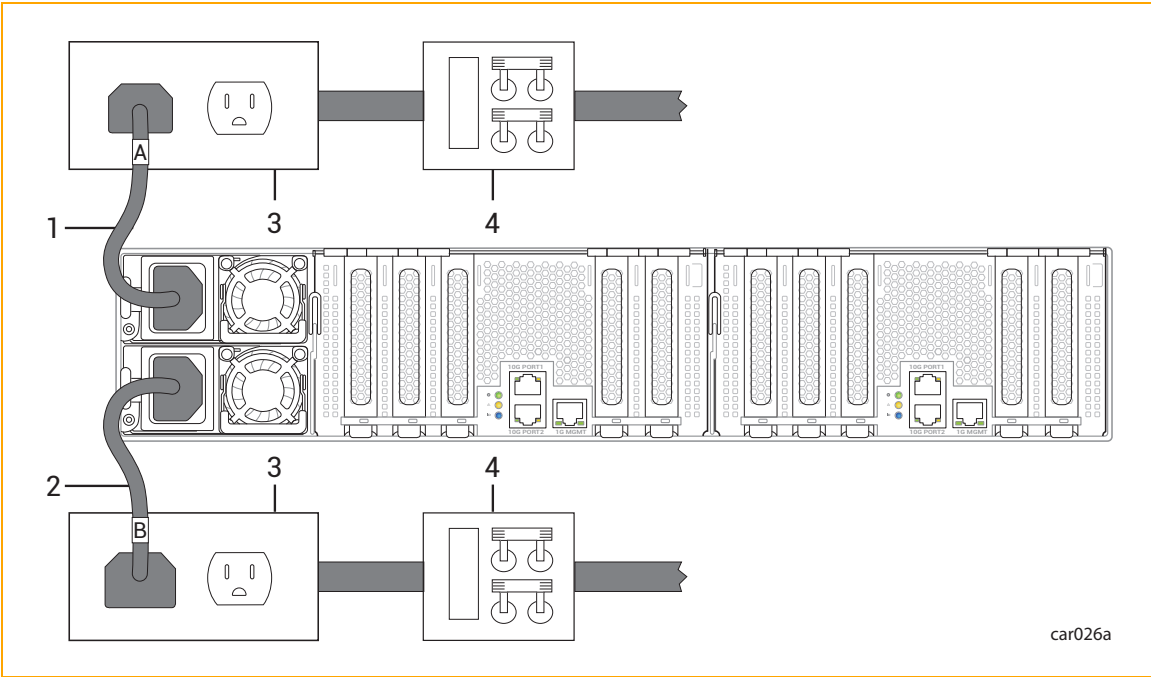
car026a

1	A-side power cord	3	AC power outlets
2	B-side power cord	4	AC power (mains) distribution circuit breakers (maximum of 20A)

To fasten the power-cord retention strap

1. Locate the power-cord retention strap on PSU A, the upper PSU at the rear of the system, as shown in Figure 3-13.

Figure 3-13: Locating the Power-Cord Retention Strap

car035a
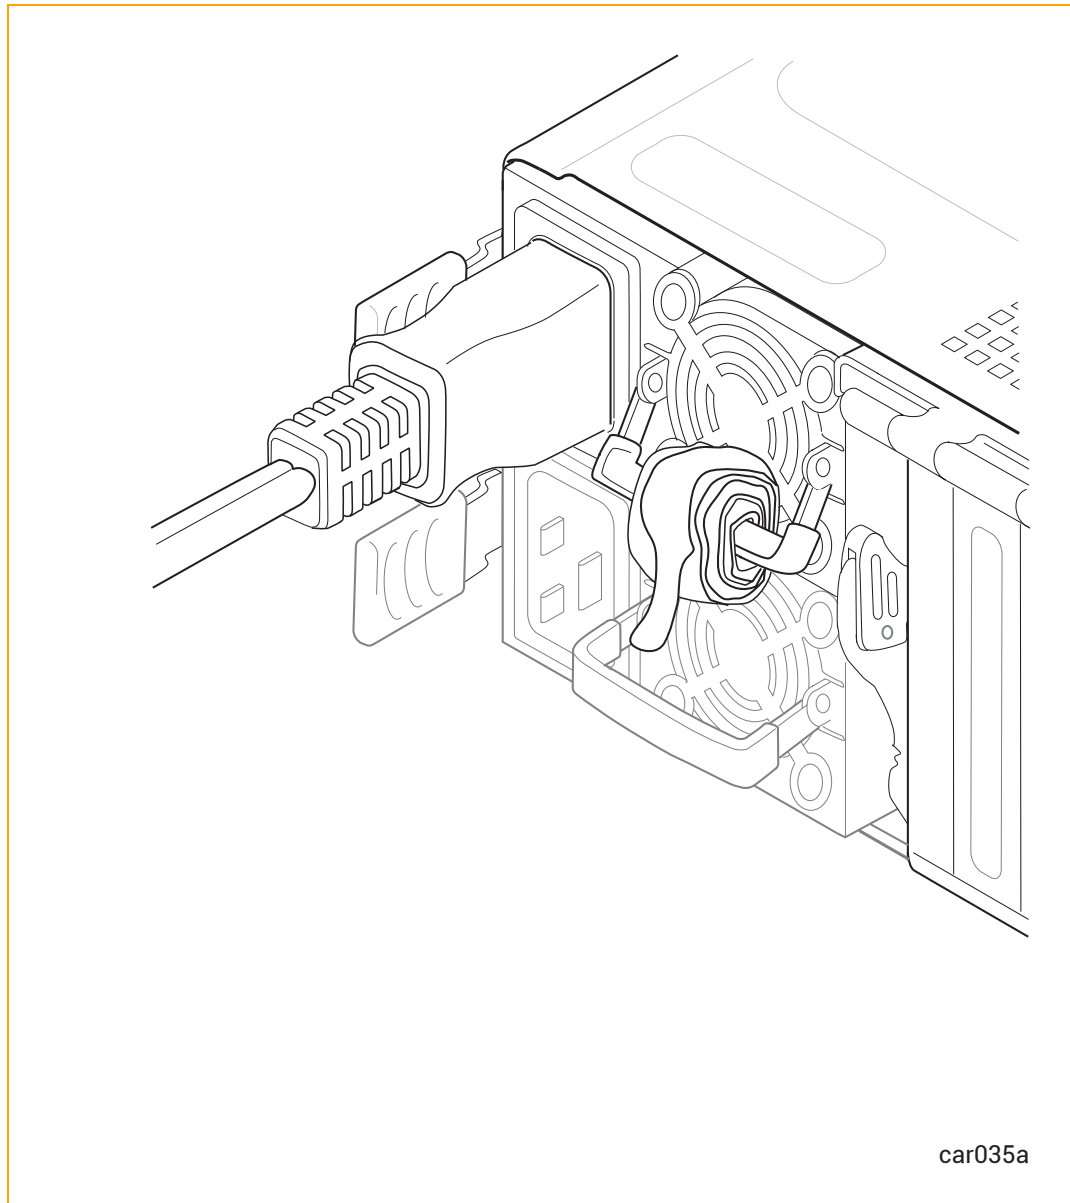
2. Pull on the end of the coiled strap to extend it, as shown in Figure 3-14.

Figure 3-14: Extending the Power-Cord Retention Strap

car035b
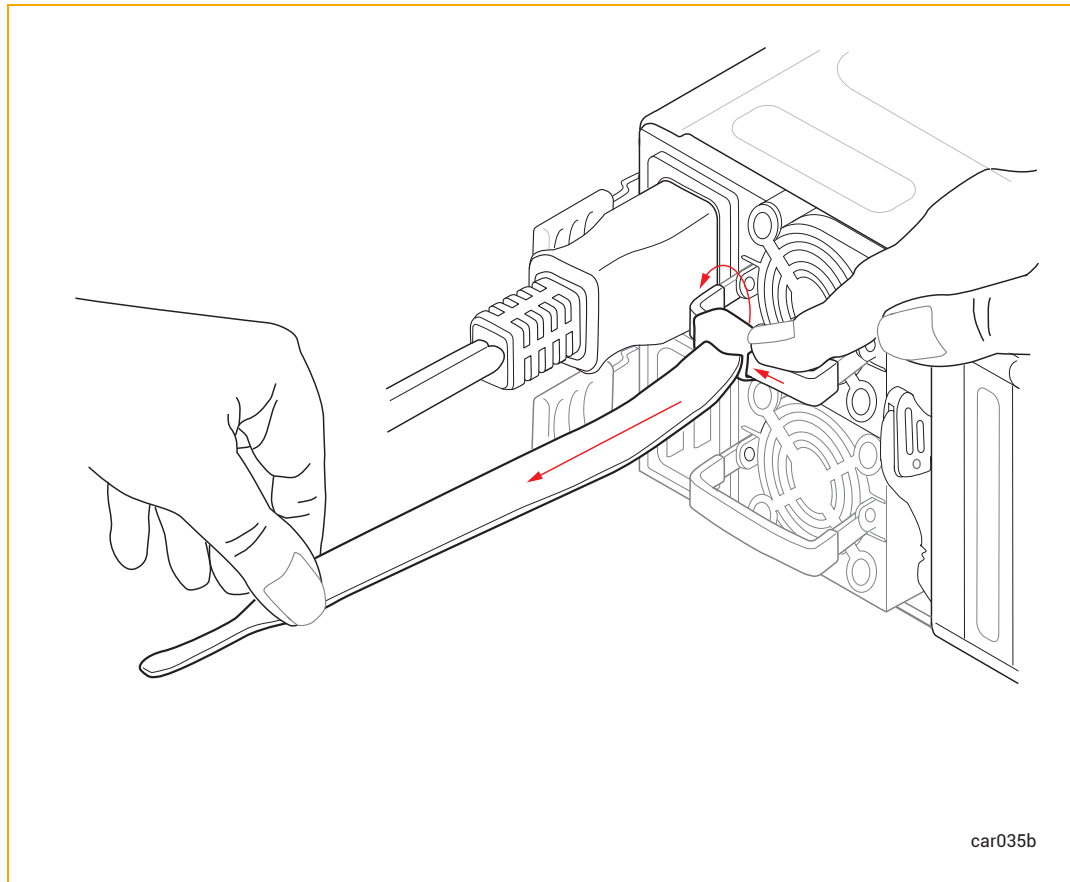
3. Wrap the strap tightly against the molded rib of the power cord to prevent the power cord from moving, as shown in Figure 3-15.

Figure 3-15: Installing the Power-Cord Retention Strap

car035c
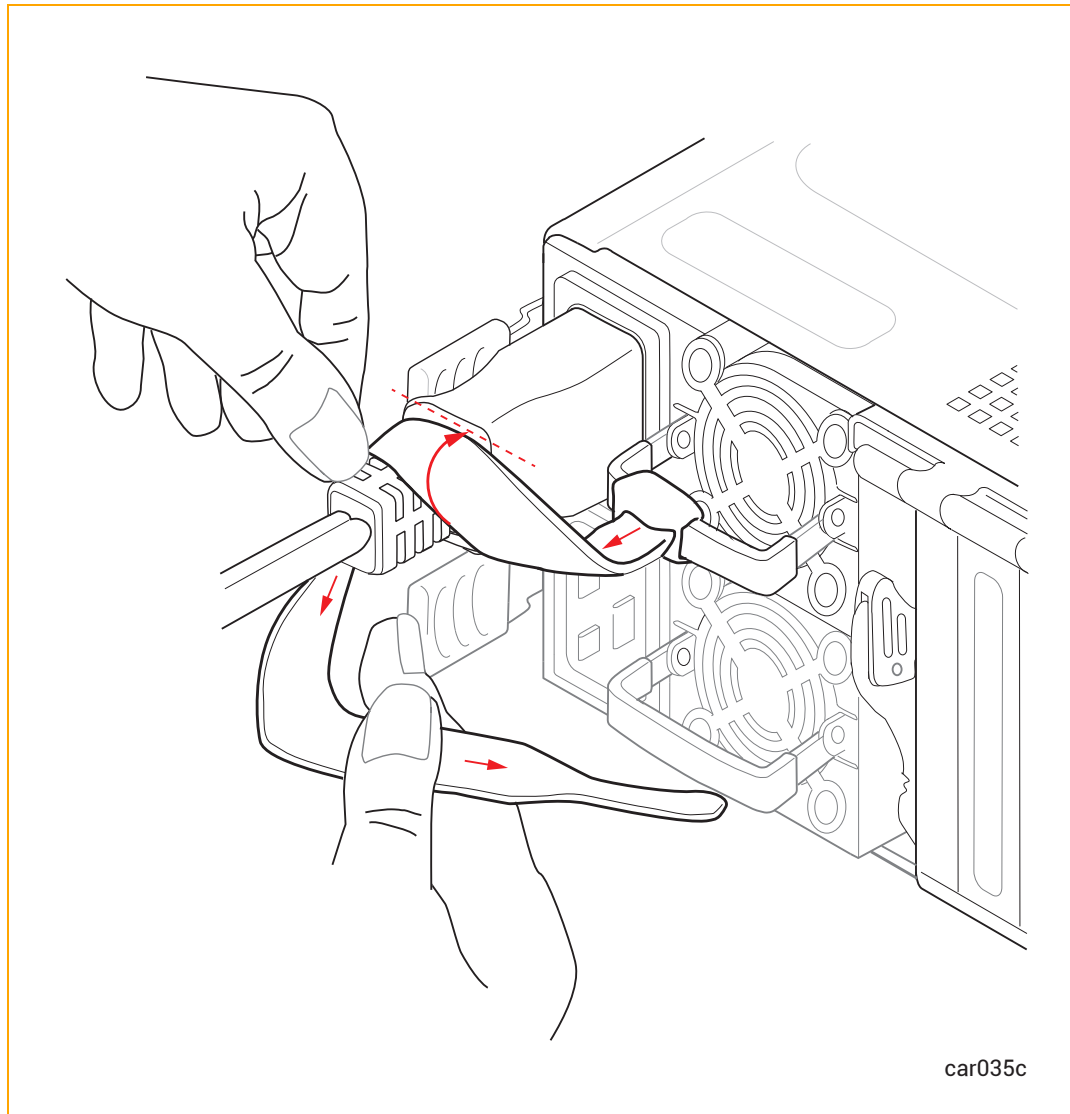
4. Continue wrapping the strap around the power cord until the strap is tightly in place, as shown in Figure 3-16.

Figure 3-16: Correct Installation of the Power-Cord Retention Strap

car035d
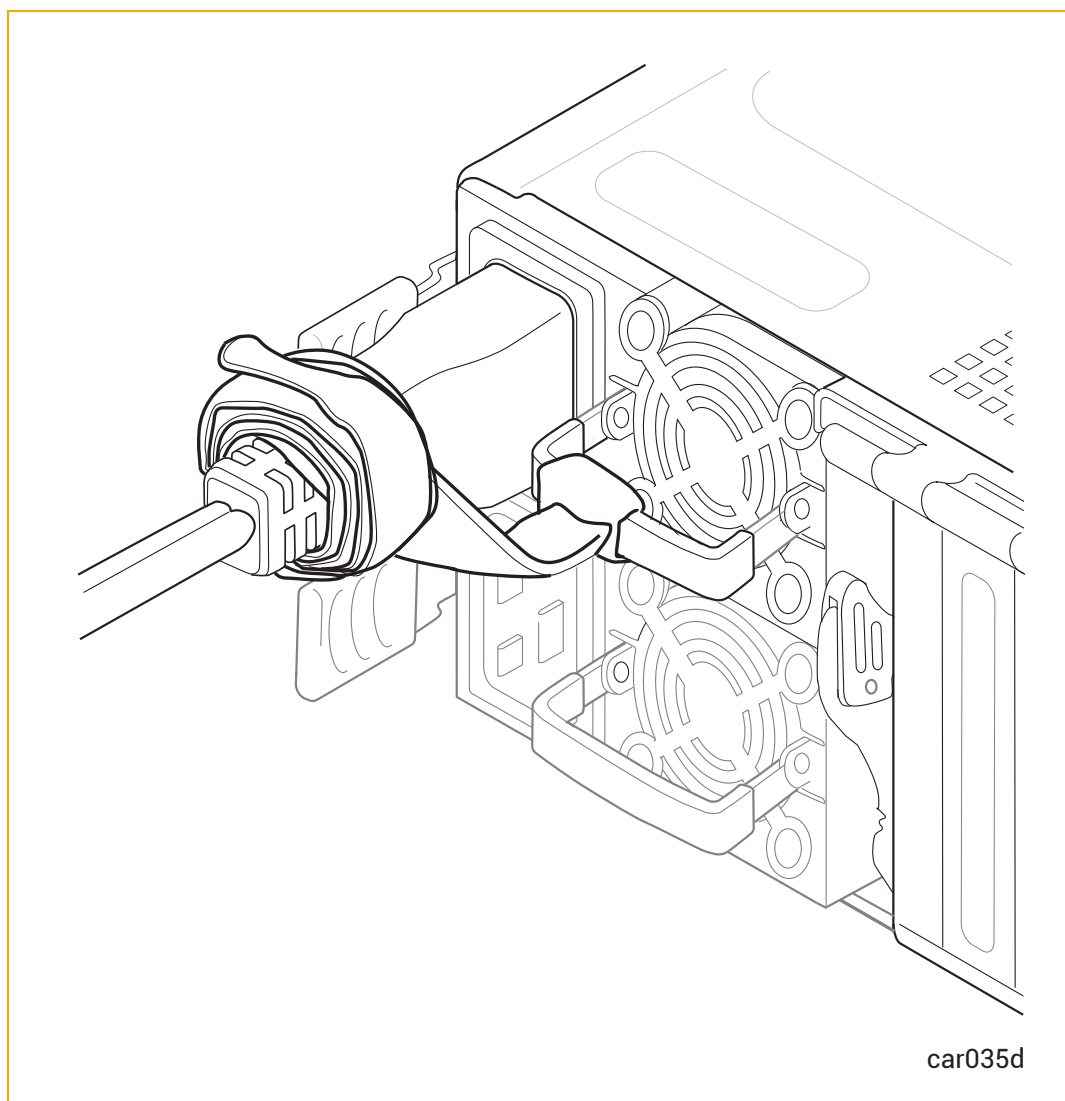
5. Repeat this procedure on PSU B, the lower PSU.

IV. Insert Blank Disks into Storage Modules

Insert a blank disk into the top-left slot of each storage module to install as system disks. (See "Disk Drives" on page 27 if you need additional information.) These disks should have the same model, firmware level, and geometry.

> **Notice**: At this time, do not install any disks that contain data in order to prevent possible problems.

Next Steps

Install the system software, as described in "Installing the System Software" on page 143.

4

Chapter 4: Operating the ztC Endurance System Hardware

To learn more about operating and troubleshooting the ztC Endurance hardware, see:

- "Hardware Operation" on page 85
- "Troubleshooting the Hardware" on page 89

Hardware Operation

To learn more about operating the ztC Endurance hardware, see:

- "Standby Power" on page 85
- "System Power" on page 86
- "Fan Speed" on page 89

Related Topics

- "Troubleshooting the Hardware" on page 89

Standby Power

When one or both of the PSU power cords are connected to live power outlets, low-level standby power is present in the system. Standby power enables the primary Baseboard Management Controller (BMC) to monitor the status of system components, even when the system power is off.

> **Note**: When you apply external power to the system, the BMC takes approximately 2 minutes to initialize. During this time, standby power, status LEDs, and the remote BMC cannot operate. The system power button may be able to operate before the BMC is fully initialized.

Standby power also allows for continuous operation of the ports used for the BMC: the 1G management ports (see "System Ports" on page 19). Therefore, while the system power is off, administrators (or, if authorized, Penguin Solutions Customer Service) can still connect to the system through the 1G management ports to diagnose problems, or to turn system power on and off from a remote location.

> ⬣ **Warning**: Standby power remains present even after the system has been powered down. To remove power to the standby devices, remove the power cords.

> ⬣ **Avertissement**: Une alimentation de réserve reste présente même après l'extinction du système. Pour couper l'alimentation des appareils de réserve, enlever les cordons d'alimentation.

Related Topics

- "System Power" on page 86

- "System Ports" on page 19

- "PSUs" on page 34

System Power

ztC Endurance systems have one power button located on the control-panel ear. Press this button to start the system.

Pressing the Power Button

The system power button functions as follows, assuming both compute modules are in the same state:

> ℹ **Note**: The power button affects both compute modules. Therefore, if one compute module is powered off and one is booted, pressing the power button results in one compute module turning on and the other shutting down.

- If the system power is off and standby power is on, pressing the power button turns the system on and boots the operating system. See "Starting Up the System" on page 86.

- If the operating system hangs or a hardware failure occurs, pressing the power button for several seconds (a "long press") initiates a forced system shutdown. (Standby power remains on.)

Starting Up the System

Make sure the system power cords are plugged into a live power source.

To boot a system and start the operating system

1. Turn on the monitor.

2. Press the power button on the control-panel ear. See Figure 4-1.

Figure 4-1: Power Button

car034

| 1 | Power button |
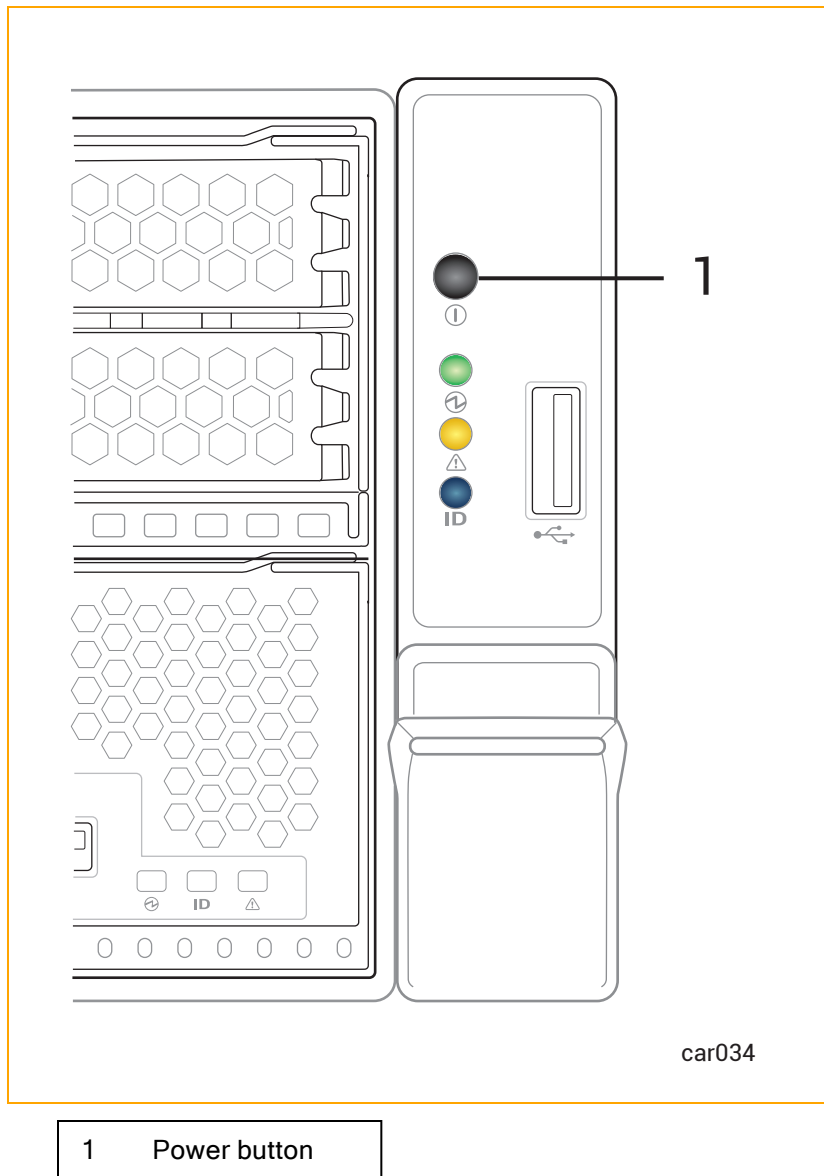
Shutting Down the System

> ⬣ **Warning**: The system's standby devices have power even after the system has been powered down. To remove power to the standby devices, you must remove the power cords.

> ⚠️ **Avertissement**: Les appareils de réserve du système restent alimentés même après l'extinction du système. Pour couper l'alimentation des appareils de réserve, vous devez enlever les cordons d'alimentation.

Exit from all applications before shutting down a system.

To shut down a Linux-based system

To shut down a Linux-based system gracefully, click **Shutdown** on the **Dashboard** page of the ztC Endurance console.

> **Notice**: Do not use the BMC Web console to power off a running ztC Endurance system. See "Unsupported BMC Web Console Features" on page 275 for more information.

When you shut down a Linux-based system, note that the following behavior is normal:

- The system stays powered on for several minutes, as all services are stopped and cached data is written to disk.

- The Ethernet activity and connection-speed LEDs for embedded Ethernet ports may remain lit even when the system is powered down, if Wake On LAN is enabled in the system BIOS.

> **Note**: Wake On LAN is supported only for the host's or BMC's connection to the 1G management network.

Enabling and Disabling Wake on LAN

The Wake On LAN (WOL) networking standard allows a computer to be turned on or awakened from sleep mode remotely via a network message.

The following instructions allow you to enable or disable WOL for ports on your ztC Endurance system.

> **Notice**: Do not change any other BIOS settings unless Penguin Solutions Customer Service has advised you to do so.

> **Note**: If you plan to enable WOL, be sure to record compute module A's original MAC addresses, which you can find in the BIOS setup utility or on the ztC Endurance console. After a Smart Exchange occurs, you need to wake compute module A with its original MAC addresses, not with its cloned MAC addresses.

To enable WOL

1. As the system starts to boot, press either the `Delete` key or the `Esc` key until the system BIOS setup utility appears.

2. When the utility screen appears, use the arrow keys to go to the `Advanced` menu.

3. Select `Intel(R) i210 Gigabit Backplane Connection` and press `Enter`.

4. Select `NIC Configuration` and then select the port that you want to enable. Press `Enter`.

5. Use the +/= keys to change `Wake On LAN` to `Enabled` and press `Enter`.

6. Press `F4` to save your changes and exit the utility.

To disable WOL

1. Download the Intel® Ethernet Flash Firmware Utility (BootUtil) from the following Web site:

   https://www.intel.com/content/www/us/en/support/articles/000005790/software/manageability-products.html

2. Use the `-WOLDISABLE` or `-WOLD` option of BootUtil to disable WOL functionality on the selected port, as described in the preceding Web site.

Fan Speed

You might notice fan-speed increases in high-temperature environments, during high system activity periods, while loading the operating system software, shutting down the system, or restarting the operating system. The increase in fan speed ensures adequate cooling under these conditions and is normal behavior.

Related Topics

- "Troubleshooting the Hardware" on page 89

Troubleshooting the Hardware

To learn more about troubleshooting the ztC Endurance hardware, see:

- "Status LEDs" on page 90

- "Internal Disk Drive Problems and Solutions" on page 99

- "Using the NMI Button" on page 100

Related Topics

- "Hardware Operation" on page 85

Status LEDs

Status LEDs are located on the control-panel ear as well as the front of each CRU module. This topic illustrates and describes these LEDs. See Figure 4-2 and Figure 4-3, as well as Table 4-1, Table 4-2, Table 4-3, Table 4-4, and Table 4-5.

You can use the ATTN LEDs to diagnose faults, as described in "Using the ATTN LEDs to Diagnose Faults " on page 97.

Figure 4-2: Front LEDs: Storage Module, Control-panel Ear, and Compute Module
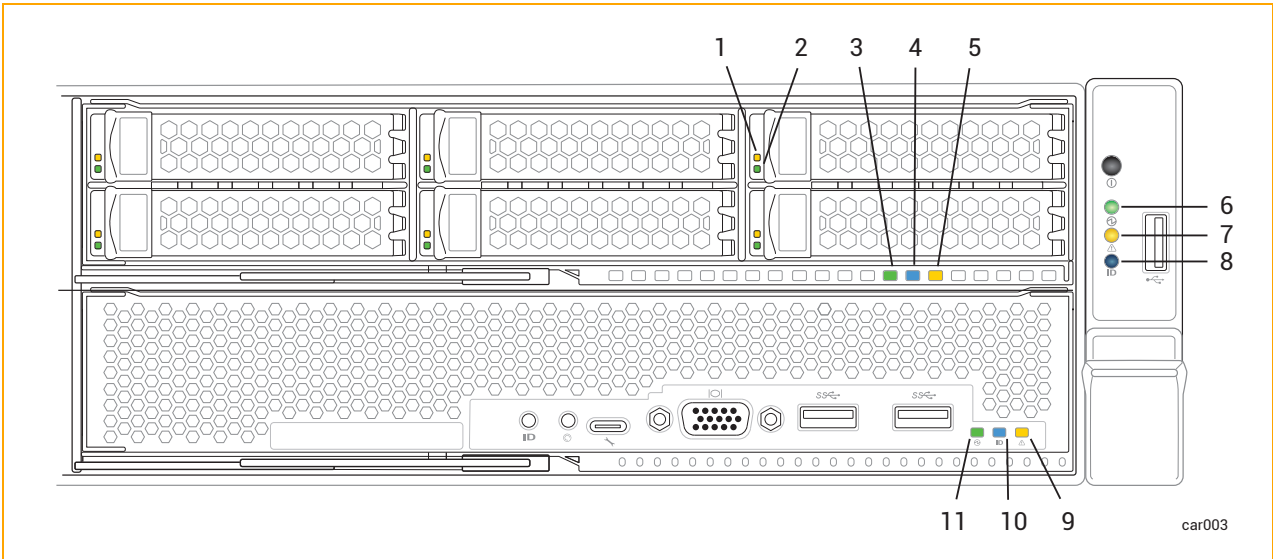
Table 4-1: Front LEDs: Storage Module

Callout	LED	State	Description
1	Fault	Off	Normal operation.
		Solid amber	The disk is broken and safe to remove. See "Removing and Inserting a Disk Drive" on page 113.
2	ACT	Off	For 6.4 TB disks, indicates that the disk is present but idle.

Table 4-1: Front LEDs: Storage Module

Callout	LED	State	Description
		Solid green	For 1.6 TB or 3.2 TB disks, indicates that the disk is present but idle.
		Blinking green	Data is being written to or read from the disk.
3	PWR	Off	DC power and standby power to the module are off.
		Solid green	DC power to the module is on.
		Blinking green	Standby power to the system is on; DC power to the module is off.
4	UID	Off	No request to identify module activated (normal operation). To identify, see "Identifying a System or Component" on page 234.
		Solid blue	Request to identify module activated; LED illuminates with DC power or with standby power.
5	ATTN	Off	Indicates one of two scenarios: • The module does not need to be replaced. • The module needs to be replaced, but the module is not safe to remove. To determine which scenario is true, check the control-panel ear's ATTN LED. If that LED is not illuminated, no module in the system needs to be replaced.
		Blinking amber	Fault; safe to replace the module. LED illuminates with DC power or with standby power. To replace, see "Removing and Replacing a Storage Module" on page 111.

Table 4-2: Front LEDs: Control-Panel Ear

Callout	LED	State	Description
6	PWR	Off	DC power and standby power to the module are off. To turn power on or off, press the power button on the control-panel ear or use the BMC. See "System Power" on page 86.
		Solid green	DC power to any module is on.
		Blinking green	Standby power is on; no module has DC power.
7	ATTN	Off	Indicates one of two scenarios: • Standby power is off. • Standby power is on, but no compute module is inserted.
		Solid green	Indicates one of two scenarios: • Standby power is on, at least one compute module is inserted, but the Stratus Management Service is not running. • None of the three scenarios that cause the LED to blink amber has occurred.
		Blinking amber	Indicates one of three scenarios: • The Linux OS is booted, the Stratus Management Service is running, and one or more modules are not inserted or need replacement. • One member of a bonded PCIe adapter or Fibre Channel cable has been removed. • A disk drive is missing from a mirrored pair, either because of physical

Table 4-2: Front LEDs: Control-Panel Ear

Callout	LED	State	Description
			removal or because it needs replacement. Check LEDs of all modules to identify the module with a fault. For a list of modules, see "System CRU Modules" on page 20. **Note**: After a total system shutdown, the ATTN LED remains in the last state it was in until you remove and restore standby power.
8	UID	Off	No request to identify module activated (normal operation). To identify, see "Identifying a System or Component" on page 234.
		Solid blue	Request to identify module activated; LED illuminates with DC power or with standby power.

Table 4-3: Front LEDs: Compute Module

Callout	LED	State	Description
9	ATTN	Off	Indicates one of two scenarios: - The module does not need to be replaced. - The module needs to be replaced, but the module is not safe to remove. To determine which scenario is true, check the control-panel ear's ATTN LED. If that LED is not illuminated, no module in the system needs to be replaced.

Table 4-3: Front LEDs: Compute Module

Callout	LED	State	Description
		Blinking amber	Fault; safe to replace the module. LED illuminates with DC power or with standby power. See "Removing and Replacing a Compute Module" on page 106.
10	UID	Off	No request to identify module activated (normal operation). To identify, see "Identifying a System or Component" on page 234.
		Solid blue	Request to identify module activated; LED illuminates with DC power or with standby power.
11	PWR	Off	The module's DC power is off, and the other compute module's DC power is on.
		Solid green	The module has DC power (standby power is also on).
		Blinking green	The system has standby power only.

Figure 4-3: Rear LEDs: PSUs and I/O Module
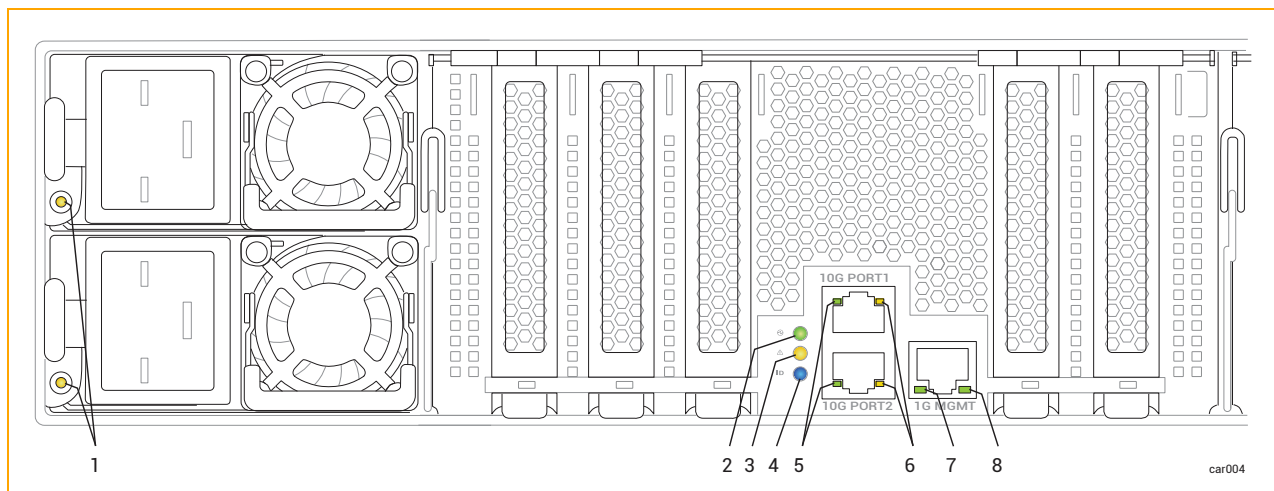
Table 4-4: Rear LEDs: PSUs

Callout	State	Description
1	Off	No AC power to any PSU.
	Solid green	Output power is on.
	Blinking green (1Hz)	The PSU is in a standby state: AC power is present at only 12Vsb.
	Blinking green (2Hz)	PSU firmware is updating.
	Solid amber	Indicates one of two scenarios: • The PSU has lost AC power (the AC cord might be unplugged); a second PSU still has AC power. • A PSU critical event is causing a shutdown because of a failure, overcurrent, short circuit, over-voltage, fan failure, over-temperature, or some other reason.
	Blinking amber (1Hz)	The PSU is still operating, though it is experiencing a warning event such as a high temperature, high power, high current, or a slow fan.

Table 4-5: Rear LEDs: I/O Module

Callout	LED	State	Description
2	PWR	Off	DC power and standby power to the module are off.
		Solid green	DC power to the module is on.
		Blinking green	Standby power to the system is on; DC power to the module is off.
3	ATTN	Off	Indicates one of two

Table 4-5: Rear LEDs: I/O Module

Callout	LED	State	Description
			scenarios: - The module does not need to be replaced. - The module needs to be replaced, but the module is not safe to remove. To determine which scenario is true, check the control-panel ear's ATTN LED. If that LED is not illuminated, no module in the system needs to be replaced.
		Blinking amber	Fault; safe to replace the module. LED illuminates with DC power or with standby power. To replace, see "Removing and Replacing an I/O Module" on page 118.
4	UID	Off	No request to identify module activated (normal operation). To identify, see "Identifying a System or Component" on page 234.
		Solid blue	Request to identify module activated; LED illuminates with DC power or

Table 4-5: Rear LEDs: I/O Module

Callout	LED	State	Description
			with standby power.
5	10 Gbps Link /ACT	Off	No link is present.
		Solid green	Link is stable.
		Blinking green	Data is being transferred over the link.
6	10 Gbps Speed	Off	No link is present.
		Solid green	Connection speed is 10 Gbps.
		Solid amber	Connection speed is less than 10 Gbps.
7	1 Gbps Link /ACT	Off	No link is present.
		Solid green	Link is stable.
		Blinking green	Data is being transferred over the link.
8	1 Gbps Speed	Off	No link is present, or link speed is less than 1 Gpbs.
		Solid green	Connection speed is 1 Gbps.

Using the ATTN LEDs to Diagnose Faults

The ATTN LED on the control-panel ear and on the CRU modules help you diagnose faults. A *fault* is an issue that requires module replacement, as opposed to other types of errors that occur on a module (for example, correctable errors below a certain threshold) that do not require module replacement.

The examples in Table 4-6 provide more information about using the ATTN LEDs to diagnose faults.

Table 4-6: Examples of Faults

ATTN LED Status	Explanation
Control-panel ear ATTN LED is blinking amber. For all modules, ATTN LEDs are off.	One or more of the modules needs replacement or is not inserted, but no module that needs replacement is safe to remove. For example, compute module A may need replacement, but it may still be the active compute module.
Control-panel ear ATTN LED is blinking amber. Compute module A ATTN LED is blinking amber. For all other modules, ATTN LEDs are off.	Compute module A needs replacement, and it is safe to remove. One or more of the other modules may need replacement or may not be inserted, but no module that needs replacement is safe to remove. For example, storage module A may need replacement, but it may contain a non-mirrored disk.
Control-panel ear ATTN LED is blinking amber. I/O module B ATTN LED is blinking amber. For all other modules, ATTN LEDs are off.	I/O module B needs replacement, and it is safe to remove. One or more of the other modules may need replacement or may not be inserted, but no module that needs replacement is safe to remove. For example, compute module A may need replacement, but it may still be the active compute module.

Related Topics

- "Identifying a System or Component" on page 234

- "System Front" on page 17

- "System Rear" on page 18

- "Planning a Site for ztC Endurance Systems" on page 37

- "Hardware Operation" on page 85

- "Removing and Replacing System Components" on page 101

- "Troubleshooting the Hardware" on page 89

Internal Disk Drive Problems and Solutions

Table 4-7 lists some problems that the disk drives may encounter, as well as some corrective actions.

Table 4-7: Troubleshooting Disk Drives

Problem	Action
System will not boot from system disk	Try to boot from the system disk's mirrored partner. To do so, disengage the suspected faulty system disk by pulling it one to two inches (2.5 to 5 cm) out of the slot. See "Removing and Inserting a Disk Drive" on page 113 for more information.
System does not recognize disk drive	Ensure that the disk drive was installed correctly. See "Removing and Inserting a Disk Drive" on page 113 for more information.
Slow response time	Mirror regeneration might be occurring. Try to schedule mirror regeneration for a time when the system is not heavily used. However, if the disk contains business-critical data, you should regenerate the mirrors as soon as possible. Many read or write operations might be queued to the disk drive. Errors might be occurring on the drive. Check the system event log.

Related Topics

- "Disk Drives" on page 27

- "Storage Module" on page 26

- "Troubleshooting the Hardware" on page 89

Using the NMI Button

To enable your authorized Penguin Solutions service representative to diagnose problems with your system, you might need to create a dump file of the system's memory while the system is running.

One way to create a dump file is to press the NMI button (see Figure 1-5) on the active compute module.

> **Notice**: Pressing the NMI button stops any running applications. Do not perform a dump operation if you want any applications to continue operating. Instead, schedule downtime for this procedure.

You may need to press the NMI button more than once to generate the dump signal. However, do not press the NMI button while a dump is in progress. This may result in an unusable dump.

Related Topics

- "System Front" on page 17

- "Compute Module" on page 21

5

Chapter 5: Removing and Replacing System Components

You can remove and replace many of the components of a ztC Endurance system. Before you remove or replace any system components, see "General Servicing Information" on page 101.

To remove and replace system components, see:

Related Topics

General Servicing Information

To perform general servicing tasks related to the unpacking, storing and repacking of system components, see:

Unpacking System Components

When you receive a system component:

- Inspect the packing slip to ensure that the correct component was received.

- Inspect the packing container for damage.

Refer to the Warranty Parts Replacement Process and Return Instructions Web Site at http://www.stratus.com/go/support/ftserver/warrantyreplace if either of the following conditions applies to your shipment:

- You received the wrong shipment.

- The package is damaged.

If there are no problems, unpack the component.

To unpack a system component

1. Observing any printed directions or warnings, carefully open the package.

2. Carefully remove the protective packaging and save it for possible reuse in storing or returning a component.

3. Verify that you received the correct component by checking the packing slip and the bar code label on the component.

4. Inspect the component for shipping damage.

5. If the component is wrapped in an antistatic bag, keep it in the bag until it is installed. If you are not going to install the component immediately, protect it from damage by repacking it in its shipping package.

Related Topics

Storing System Components

If you need to store a system component, keep it in the protective packaging and take the following precautions:

- Ensure that the component is right-side up by following the directions on the package.

- Comply with all warning labels.

- Avoid placing the component in a location where heavier items might be placed on it.

- Protect the component from exposure to dust, electromagnetic fields, vibrations, extreme heat, or any other environmental conditions that might damage the component.

- Avoid prolonged storage that could reduce the useful life of the component.

If you are uncertain about the proper storage conditions for a specific component, see "Getting Help" on page 295. Also, see "Preparing to Replace System Components" on page 105 for information about how to protect components from exposure to static electricity.

Related Topics

- "Removing and Replacing System Components" on page 101

Repacking System Components

When returning a defective or incorrect system component, package the component using the original packing materials or the replacement-component packing materials. Ensure that the component is adequately protected from the following:

- Dents, scratches, and impact damage

- Static electricity (using the pink or blue electrostatic discharge (ESD) bag)

- Accidental opening of the package in transit

- Detachment of the address label

If you are uncertain that the component will be adequately protected during shipment, see "Getting Help" on page 295.

Related Topics

- "Removing and Replacing System Components" on page 101

Replacing System Components

> **Caution**: Do not remove or service any I/O modules, storage modules, or compute modules if the ztC Endurance console is reporting any PSU faults.
>
> You must first restore full PSU functionality to both PSUs before removing any other modules for additional service (that is, the LED on both PSUs must be solid green).
>
> **In extreme circumstances, failure to follow this service sequence may result in compromised power integrity to the ztC Endurance system.**

> **Caution**: Ne retirez ni ne procédez à l'entretien d'aucun module d'E/S, de stockage ou de calcul si la console ztC Endurance signale une défaillance d'un bloc d'alimentation.
>
> Vous devez d'abord rétablir le bon fonctionnement des deux blocs d'alimentation avant de retirer tout autre module pour un entretien supplémentaire (autrement dit, les voyants DEL des deux blocs d'alimentation doivent rester allumés en vert sans clignoter).
>
> **Dans des situations extrêmes, le non-respect de cette séquence d'entretien risque de compromettre l'intégrité de l'alimentation électrique du système ztC Endurance.**

Before replacing a system component, do the following:

- Confirm that the component has failed and needs to be replaced. You can view failure information from the status LEDs on each component, from your system's event or error logs, and from viewing the **Dashboard** page of the ztC Endurance console.

- Determine if a component needs to be taken out of service.

- Observe the proper safety precautions by following the precautions listed in "General Safety Precautions" on page 105.

After replacing a component, do the following:

1. Place the component in service if it is out of service.

2. Verify system operation.

3. Package and ship the failed component. See "Repacking System Components" on page 103.

Related Topics

- "Removing and Replacing System Components" on page 101

- "Status LEDs" on page 90

- "Dashboard Page" on page 224

General Safety Precautions

- To help identify problems that might occur during the removal and replacement of system components, read the entire procedure before performing it.

- Provide enough space and light to perform the procedure safely and accurately.

- Conduct all activities in compliance with all applicable industry safety standards and practices.

- Take proper grounding precautions. When handling components, especially those with exposed integrated circuits, use proper grounding procedures to avoid ESD damage to the component. See the information about using ESD precautions in "Safety Precautions" on page 51.

- When performing maintenance procedures, do not wear conductive articles or material such as rings, bracelets, keys, chains, garments with metallic thread, and other similar items.

Related Topics

- "Safety Precautions" on page 51

Preparing to Replace System Components

No tools are required to replace customer-replaceable components in a ztC Endurance system.

Working on the physical system can expose electronic components to potential damage from electrostatic discharge (ESD). Whenever working on a system:

- Use ESD precautions.

- Observe all notes, notices, cautions, and warnings.

Using ESD Precautions

> **Notices**: To avoid damaging ESD-sensitive components during handling, always take the following precautions.
>
> - Ground yourself before working with system components. Put a grounding strap on your wrist and attach its other end to some suitable grounding point, such as a computer-system cabinet.
>
> - Discharge static electricity by touching an unpainted portion of the system just before handling ESD-sensitive parts.
>
> - Store PCIe adapters in their static-protective envelope until you are ready to install them in the system.
>
> - Hold a PCIe adapter by its edges.

Observing Cautions and Warnings

Observe the following warnings, cautions, notices, and notes.

> **Warning**: If a component is not customer-replaceable, do not attempt to remove or replace it. Doing so may result in serious personal injury and/or damage to the system.

> **Avertissement**: Si un composant n'est pas remplaçable par le client, n'essayez pas de le retirer ou de le remplacer. Cela pourrait provoquer des blessures graves et/ou endommager le système.

> **Notices**: When replacing a system component, always observe the following guidelines.
>
> - Remove a component only when its status LEDS indicate that it is safe to remove. See "Status LEDs" on page 90 for information about each component's status LEDs.
>
> - When necessary, see "Status LEDs" on page 90 to learn how to identify a failed component and take it offline before servicing it.
>
> - Read the entire procedure so that you fully understand it. If any part of the procedure is not clear, see "Getting Help" on page 295.

Related Topics

- "Safety Precautions" on page 51

- "General Safety Precautions" on page 105

Removing and Replacing a Compute Module

> **Cautions**:
>
> 1. Use ESD precautions when performing these procedures.
>
> 2. Read "Replacing System Components" on page 104 before you remove any system components.

> **Précautions**:
>
> 1. Utilisez les précautions ESD lors de l'exécution de ces procédures.
>
> 2. Lisez "Replacing System Components" on page 104 avant de retirer des composants du système.

To remove a compute module

1. Place the compute module into maintenance mode as described in "Using Maintenance Mode" on page 235.

2. After the compute module is in maintenance mode, shut it down as described in "Shutting Down and Restarting a Compute Module" on page 237.

3. Disconnect all cables from the compute module.

4. Use your index finger to unlatch the release lever on the compute module.

5. Grasp the release lever and pull it away from the system chassis until the module extends an inch or two from the system chassis. See Figure 5-1.

Figure 5-1: Removing Compute Module
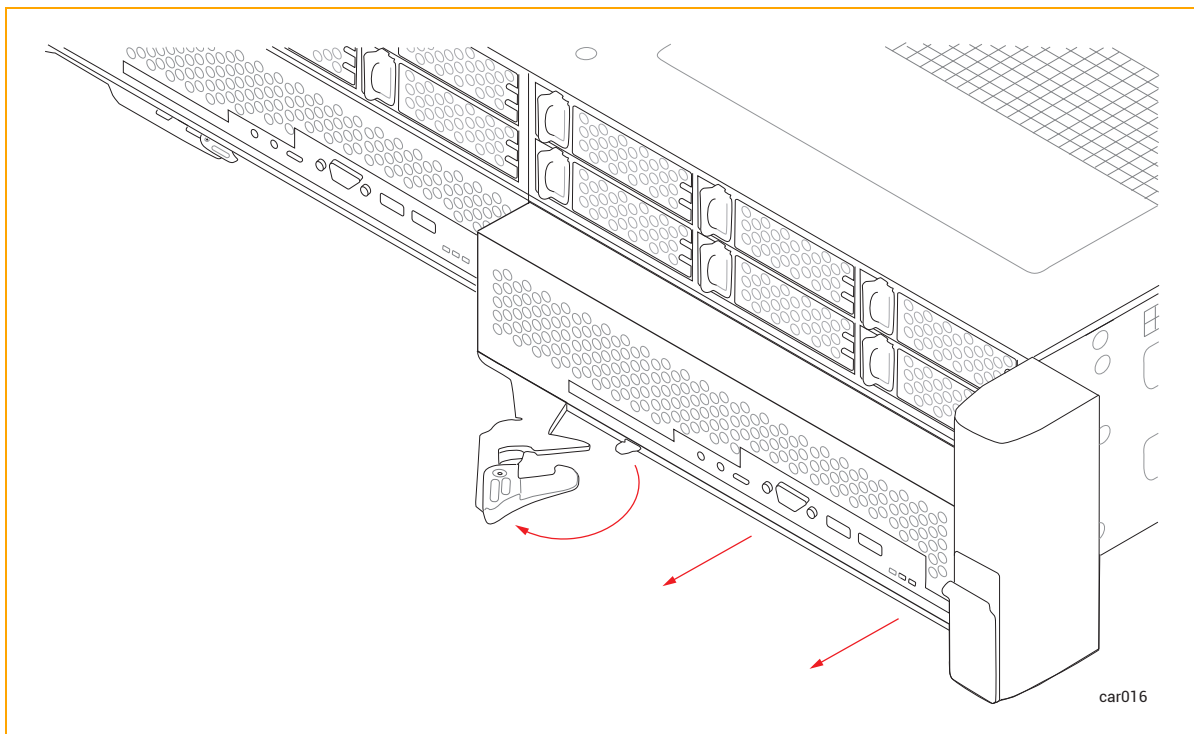
6. Holding the module on each side, pull it straight out of the system chassis and place it on a flat, stable surface.

To replace a compute module

> **Caution**: Make sure that the left compute module is fully inserted and latched before you remove the right compute module. Otherwise, damage can result, and the right compute module cannot be removed.

> **Précaution**: Assurez-vous que le module de calcul de gauche est complètement inséré et verrouillé avant de retirer le module de calcul de droite. Sinon, des dommages peuvent en résulter et le module de calcul de droite ne peut pas être retiré.

1. Check the "Hardware Page" on page 229 of the ztC Endurance console to make sure that the compute module is in maintenance mode.

2. Extend the release lever on the compute module to the left until it is fully open.

3. Holding the compute module on each side, slide it into one of the bottom slots until it stops. Ensure that the module is fully seated and aligns with the system chassis.

4. Push on the center of the compute module with one hand while simultaneously using your other hand to push the release lever toward the system chassis until it clicks. See Figure 5-2.

Figure 5-2: Replacing Compute Module

car048
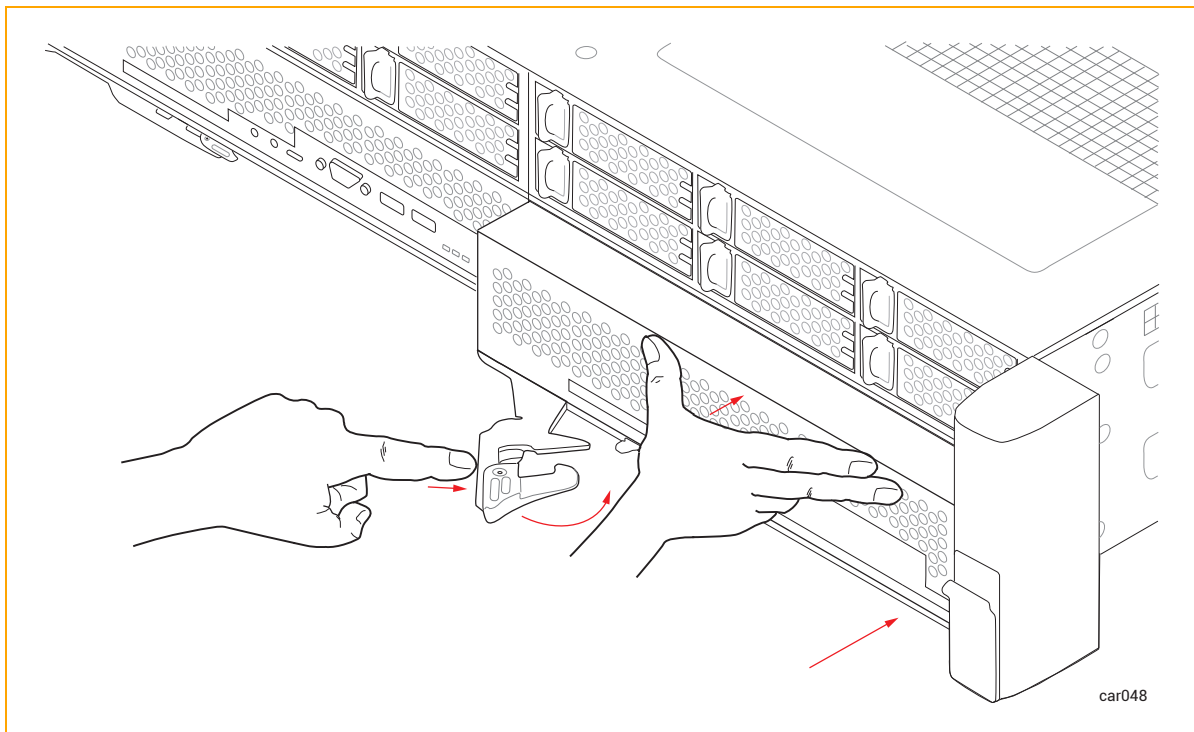
5. Push on the center of the compute module once more to make sure it is fully seated in the system chassis.

6. Connect all of the cables.

7. Remove the compute module from maintenance mode as described in "Using Maintenance Mode" on page 235.

> **Notice**: If you insert a replacement compute module but do not remove it from maintenance mode, the system software cannot update the BMC, BIOS, or Standby OS. The system software also posts an alert to indicate that it cannot update the standby compute module until maintenance mode is cleared.

Related Topics

- "Compute Module" on page 21

- "Memory in ztC Endurance Systems" on page 22

Replacing and Adding RDIMMs

You install RDIMMs when upgrading system memory or replacing a faulty RDIMM. Before you install RDIMMs, read "Memory in ztC Endurance Systems" on page 22.

> **Cautions**:
>
> 1. Use ESD precautions when performing these procedures.
>
> 2. Read "Replacing System Components" on page 104 before you remove any system components.

> **Précautions**:
>
> 1. Utilisez les précautions ESD lors de l'exécution de ces procédures.
>
> 2. Lisez "Replacing System Components" on page 104 avant de retirer des composants du système.

> **Note**: When you install RDIMMs, you do not need to schedule any downtime because you remove only one compute module at a time.

> **Notice**: If you require assistance while adding RDIMMs, your authorized Penguin Solutions service representative can provide onsite installation services for a fee.

To remove RDIMMs

1. Remove the compute module whose RDIMMs you are upgrading.

2. Push down on the ejector levers at each end of the RDIMM to be removed. See Figure 5-3.

Figure 5-3: Removing RDIMM
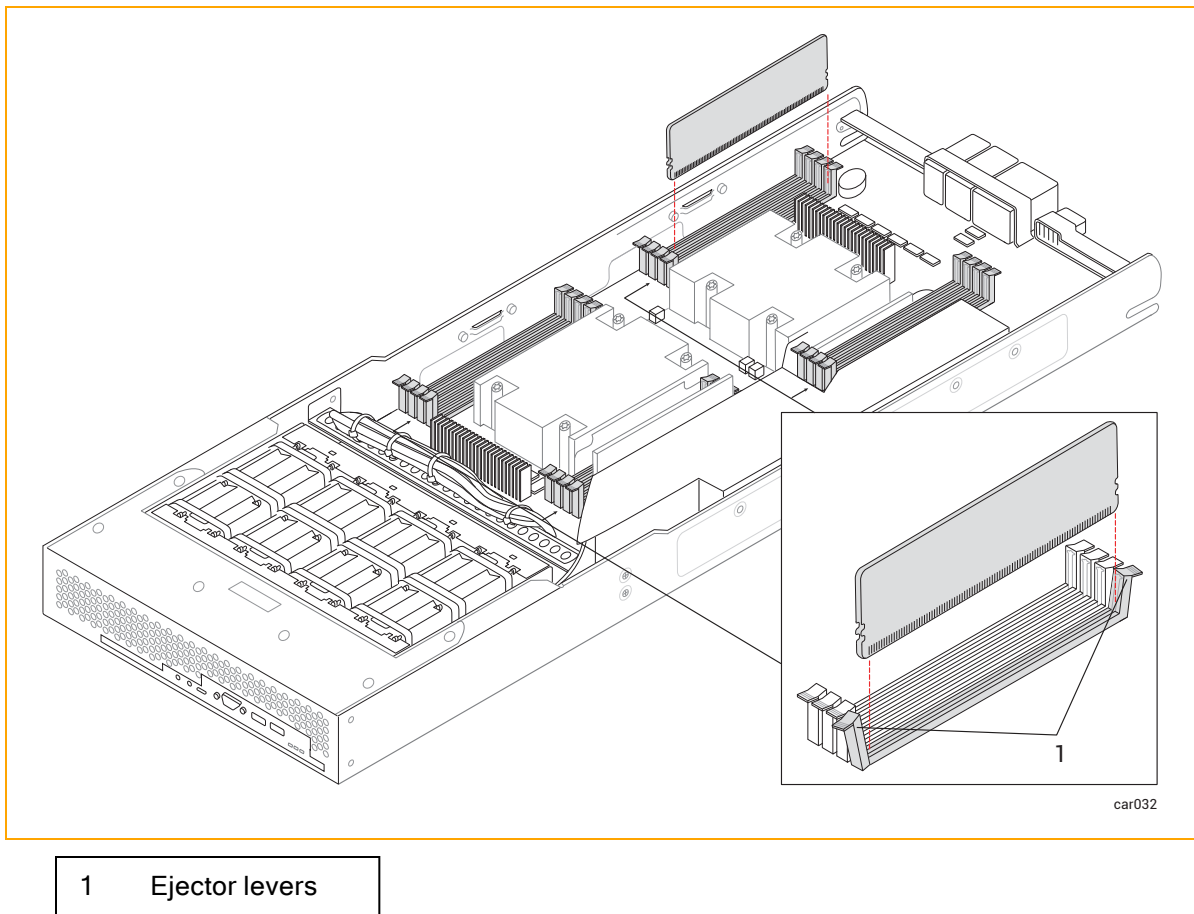
1	Ejector levers

3. Pull the RDIMM straight up to remove it from the connector.

To install RDIMMs

1. Make sure you are installing the RDIMM in the correct slot. See Figure 1-6.

2. With the RDIMM connector ejector levers in the open position (down), align the small slots on the base of the RDIMM with the notches at the base of the connector. Push the RDIMM down firmly into place. Make sure it is fully seated.

3. Ensure that the connector's ejector levers are raised and fully locked into place. If they are not, eject and reseat the RDIMM.

4. Repeat Steps 1 and 2 for any additional RDIMMs.

5. After installing the RDIMMs, replace the compute module.

Related Topics

- "Memory in ztC Endurance Systems" on page 22

- "Compute Module" on page 21

Removing and Replacing a Storage Module

Cautions:

1. Use ESD precautions when performing these procedures.

2. Read "Replacing System Components" on page 104 before you remove any system components.

Précautions:

1. Utilisez les précautions ESD lors de l'exécution de ces procédures.

2. Lisez "Replacing System Components" on page 104 avant de retirer des composants du système.

To remove a storage module

1. Use your index finger to unlatch the release lever on the storage module.

2. Grasp the release lever and pull it away from the system chassis until the storage module extends an inch or two from the system chassis. See Figure 5-4.

Figure 5-4: Removing Storage Module

3. Holding the storage module on each side, pull it straight out of the system chassis and place it on a flat, stable surface.

To replace a storage module

> **Caution**: Make sure that the left storage module is fully inserted and latched before you remove the right storage module. Otherwise, damage can result, and the right storage module cannot be removed.

> **Précaution**: Assurez-vous que le module de stockage de gauche est complètement inséré et verrouillé avant de retirer le module de stockage de droite. Sinon, des dommages peuvent en résulter et le module de stockage droit ne peut pas être retiré.

1. Extend the release lever on the storage module to the left until it is fully open.

2. Holding the storage module on each side, slide it into one of the top slots until it stops. Ensure that the module is fully seated and aligns with the system chassis.

3. Push on the center of the storage module with one hand while simultaneously using your other hand to push the release lever toward the system chassis until it clicks. See Figure 5-5.

Figure 5-5: Replacing Storage Module

car050
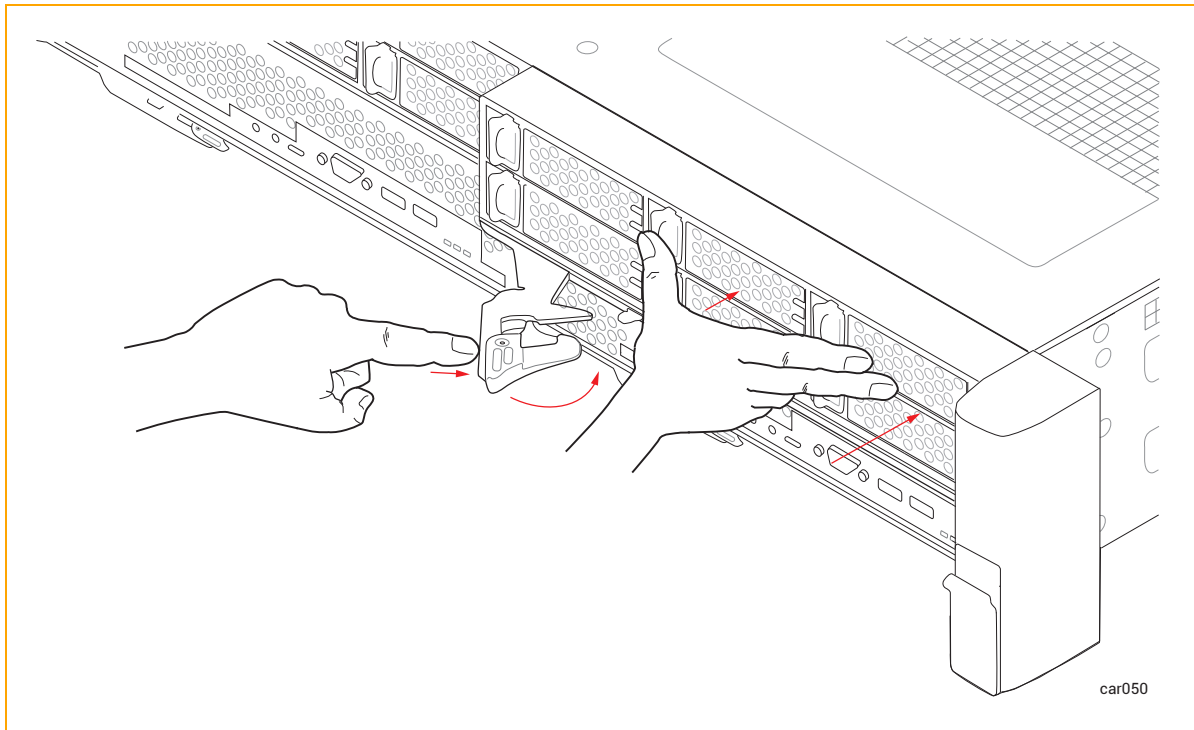
4. Push on the center of the storage module once more to make sure it is fully seated in the system chassis.

Related Topics

- "Storage Module" on page 26

- "Removing and Inserting a Disk Drive" on page 113

Removing and Inserting a Disk Drive

You can remove disk drives from a storage module to replace them or to perform other maintenance procedures. You do not need to power down or remove the storage module to perform this procedure.

> **Caution**: Use ESD precautions when performing these procedures.

> **Précaution**: Utilisez les précautions ESD lors de l'exécution de ces procédures.

To remove a disk drive

1. Slide the button on the release lever to the left until the release lever opens. See Figure 5-6.

   Figure 5-6: Removing Disk Drive
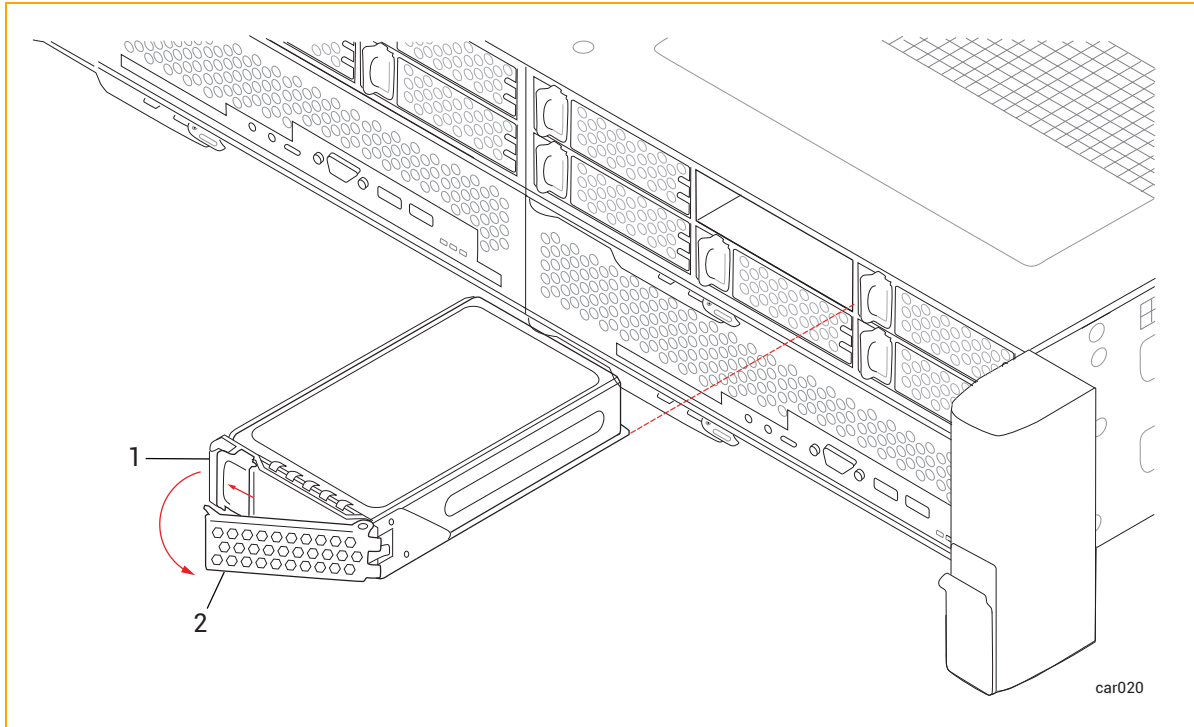
1	Release button
2	Release lever

2. Pull the release lever and remove the disk drive.

To insert a disk drive

> **Note**: You should be familiar with the information provided in "Disk Drives" on page 27 before you insert a disk drive.

1. With the release lever in the fully open position, insert the disk drive into the slot at the slight angle shown in Figure 5-7. Stop when approximately three inches of the disk drive extend from the slot.

Figure 5-7: Inserting Disk Drive: Step 1

car042c
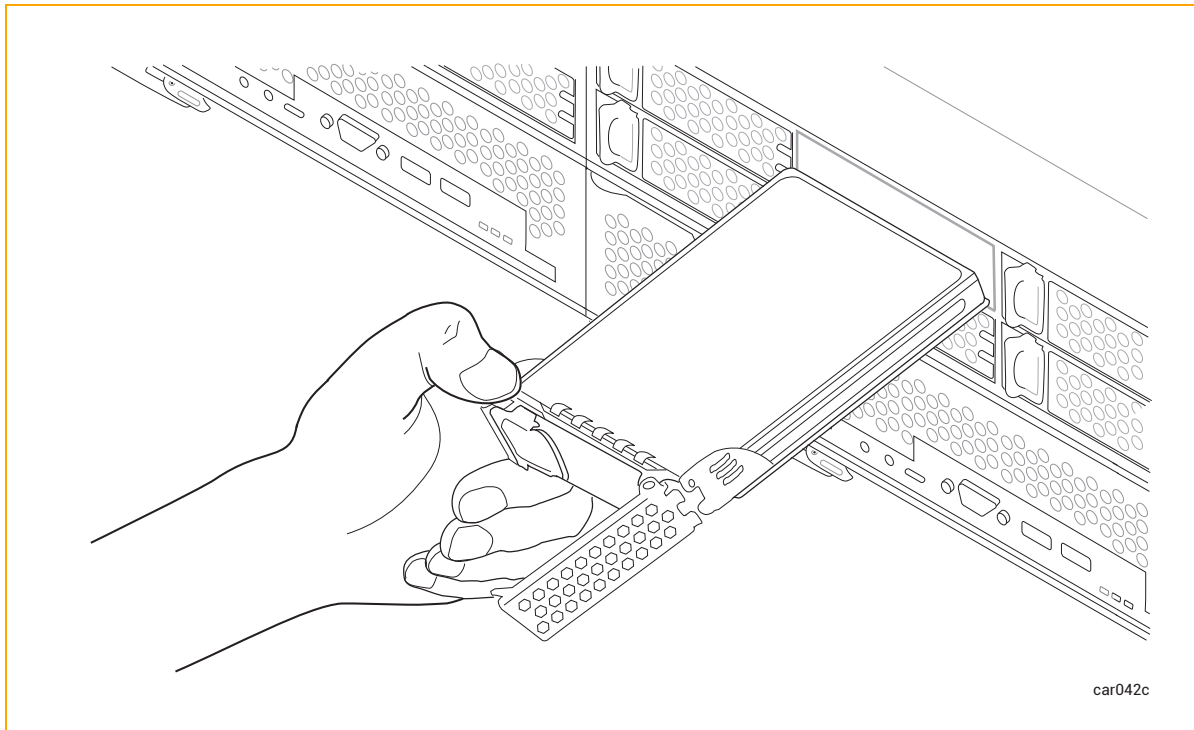
2. With the release lever still in the fully open position, use your left thumb to apply pressure on the release button while slowly inserting the disk drive into the slot. Stop when the two tabs on the right side of the release lever are flush with the slot opening. See Figure 5-8.

Figure 5-8: Inserting Disk Drive: Step 2

car043a
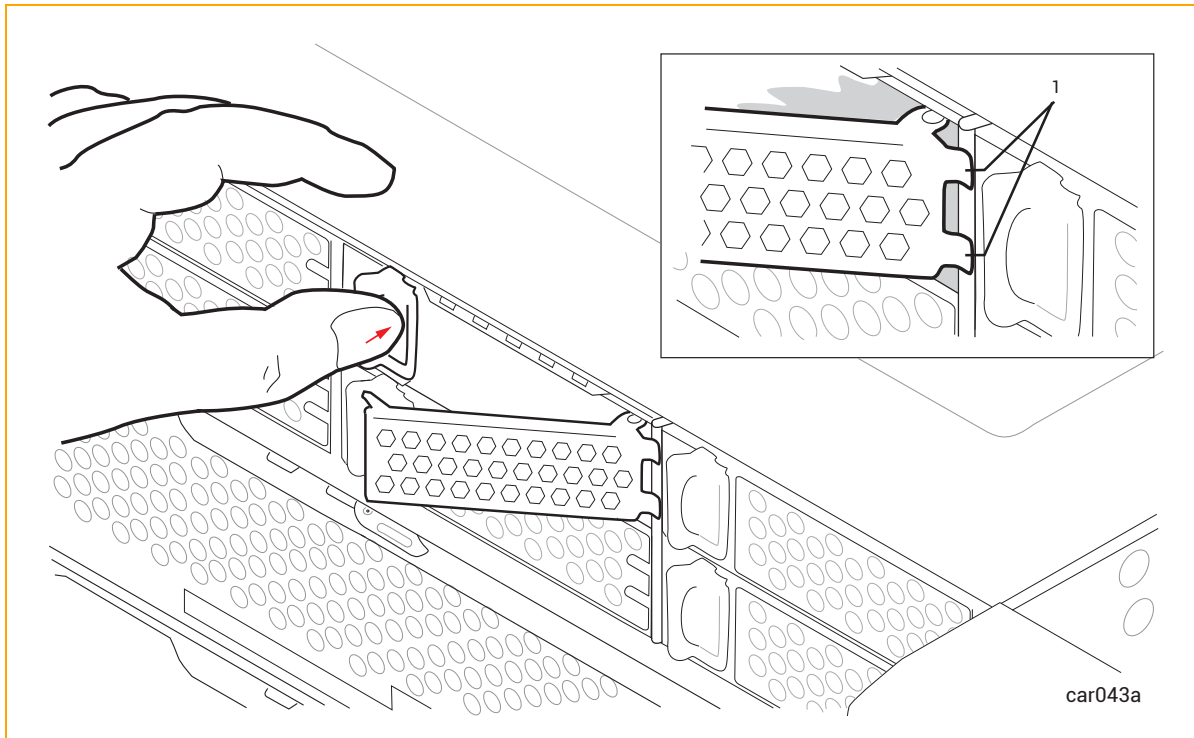
1	Tabs

3. While continuing to apply pressure to the release button with your left thumb, use your right thumb to push the release lever until it locks into place. See Figure 5-9.

Figure 5-9: Inserting Disk Drive: Step 3

car044
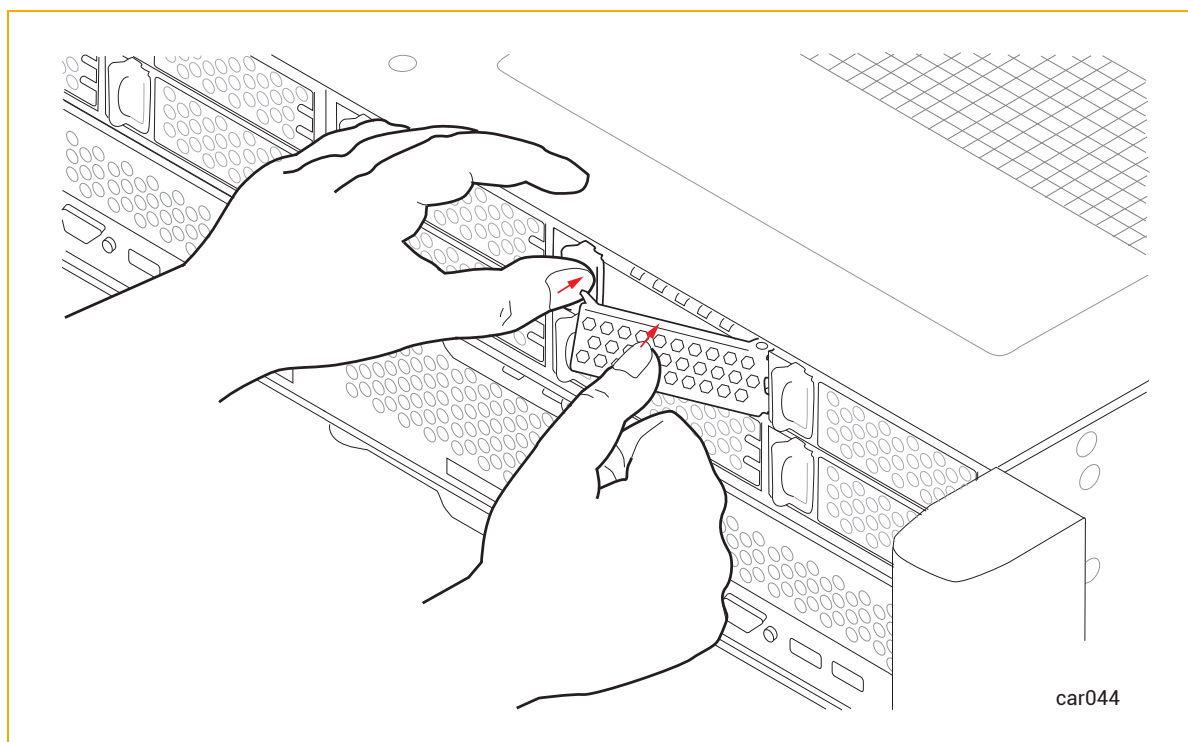
4. Push firmly again to make sure that the disk drive is fully seated. See Figure 5-10.

Figure 5-10: Inserting Disk Drive: Step 4
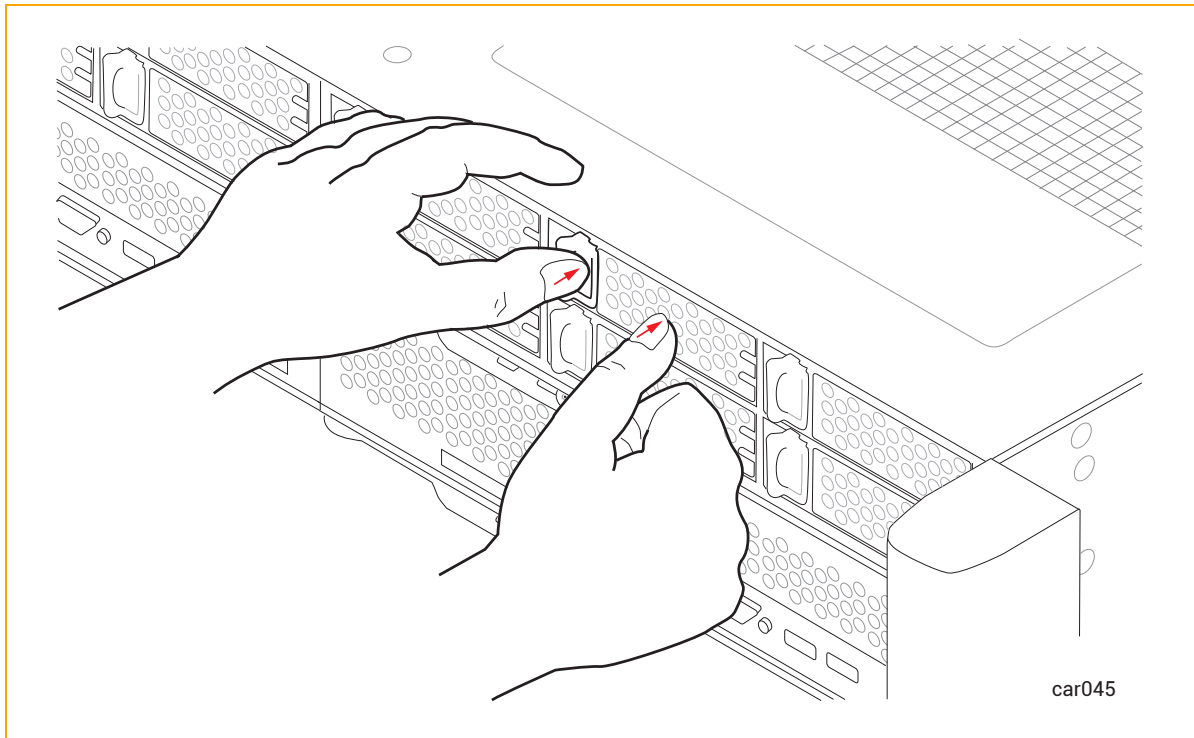
5. If the ztC Endurance system is powered on, check the disk drive's ACT LED. If the LED is solid green, the disk drive is inserted correctly. If the LED is off, the disk drive is not inserted correctly; remove the disk drive and repeat this procedure.

Related Topics

- "Disk Drives" on page 27

- "Internal Disk Drive Problems and Solutions" on page 99

- "Storage Module" on page 26

- "System Specifications" on page 289

Removing and Replacing an I/O Module

> **Cautions**:
>
> 1. Use ESD precautions when performing these procedures.
>
> 2. Read "Replacing System Components" on page 104 before you remove any system components.

Précautions:

1. Utilisez les précautions ESD lors de l'exécution de ces procédures.

2. Lisez "Replacing System Components" on page 104 avant de retirer des composants du système.

To remove an I/O module

1. Disconnect all cables from the I/O module.

2. Use your index finger to unlatch the release lever on the I/O module.

3. Grasp the release lever and pull it down until it is fully extended and the I/O module extends an inch or two from the system chassis. See Figure 5-11.

Figure 5-11: Removing I/O Module

car021
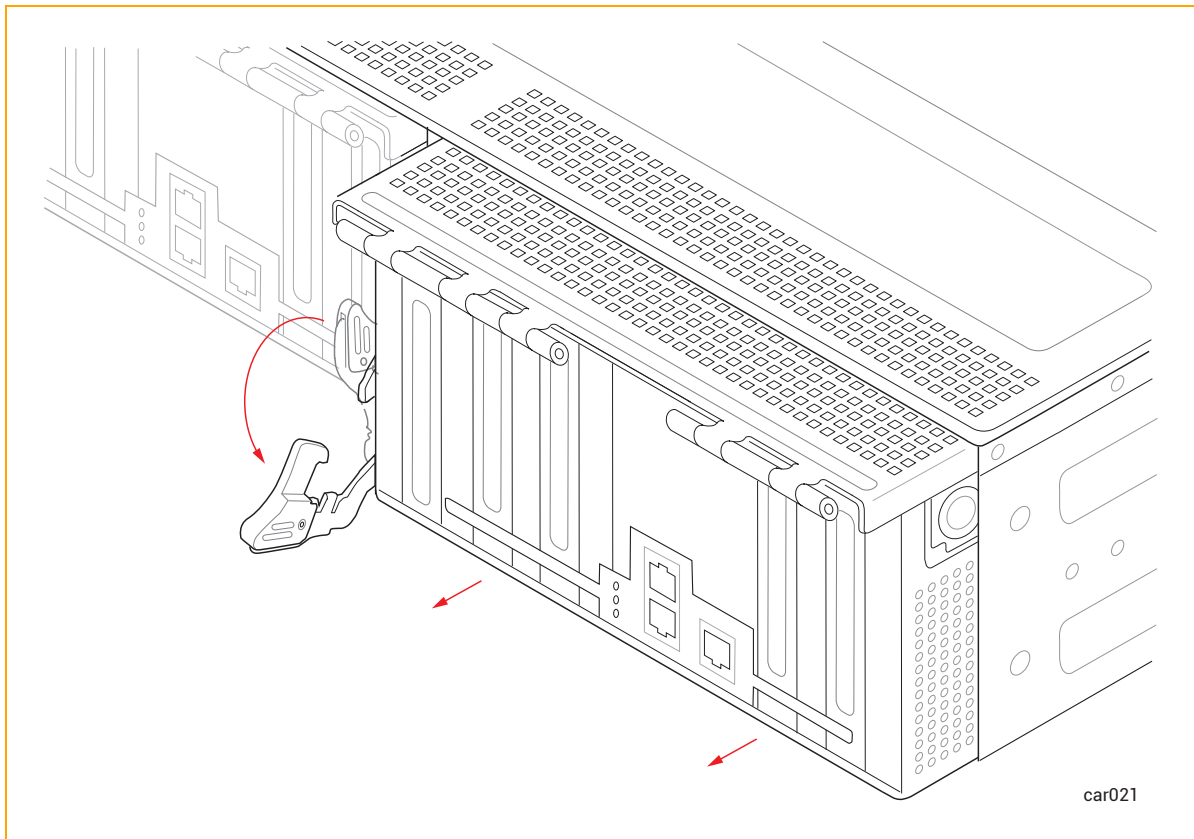
Caution: The release lever must be fully extended as shown in the figure before you perform step 4.

> **Précaution**: Le levier de déverrouillage doit être complètement étendu comme indiqué sur la figure avant d'effectuer l'étape 4.

4. Holding the I/O module on each side, pull it straight out of the system chassis and place it on a flat, stable surface.

To replace an I/O module

1. Extend the release lever on the I/O module until it is fully open.

2. Slide the I/O module into the system chassis until it stops. Ensure that the module is fully seated and aligns with the system chassis.

3. Push on the center of the I/O module with one hand while simultaneously using your other hand to push up the release lever until it clicks. See Figure 5-12.

Figure 5-12: Replacing I/O Module
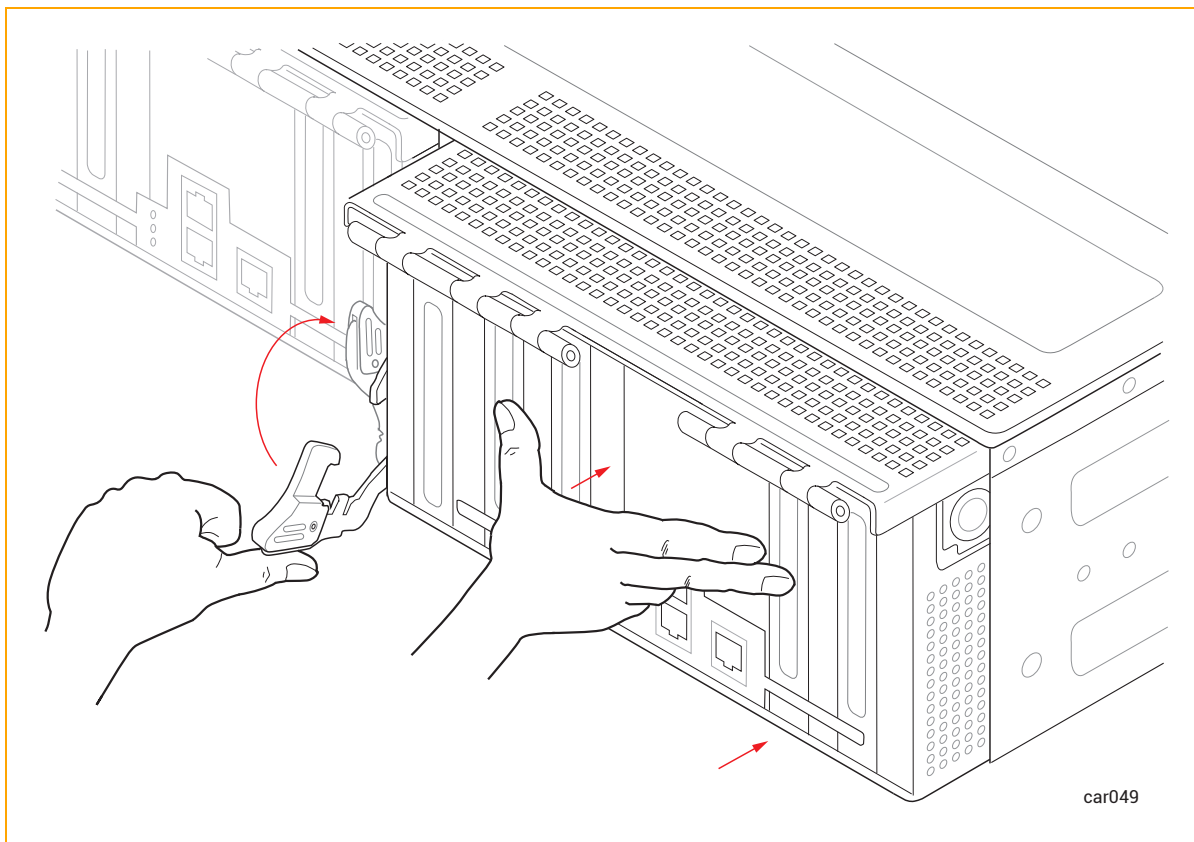
4. Push on the center of the I/O module once more to make sure it is fully seated in the system chassis.

5. Connect all of the cables.

Related Topics

- "I/O Module" on page 29

- "Removing and Installing a PCIe Adapter" on page 121

- "PCIe Adapters" on page 30

Removing and Installing a PCIe Adapter

Cautions:

1. Use ESD precautions when performing these procedures.

2. Read "Replacing System Components" on page 104 before you remove any system components.

Précautions:

1. Utilisez les précautions ESD lors de l'exécution de ces procédures.

2. Lisez "Replacing System Components" on page 104 avant de retirer des composants du système.

Notes:

1. You do not need to shut down a ztC Endurance system when you remove or install PCIe adapters. You simply remove one I/O module, remove/install PCIe adapter(s), and reinstall the I/O module. Repeat this procedure for the other I/O module.

2. Supported PCIe adapters that contain small form-factor pluggable (SFP) modules take up the length of an entire PCIe slot. To more easily remove or install these longer adapters, first remove the SFP modules as noted in the following procedures.

To remove a PCIe adapter

Notices:

1. When handling PCIe adapters, hold them securely by their edges. Do not touch their electronic elements.

2. Always store a PCIe adapter in its static-protective envelope if it is not installed in the system.

1. Remove the I/O module.

2. Disengage the metal cage by pushing the buttons on each side of the I/O module, and then flip the metal cage open. See Figure 5-13.

Figure 5-13: Opening Metal Cage

car022
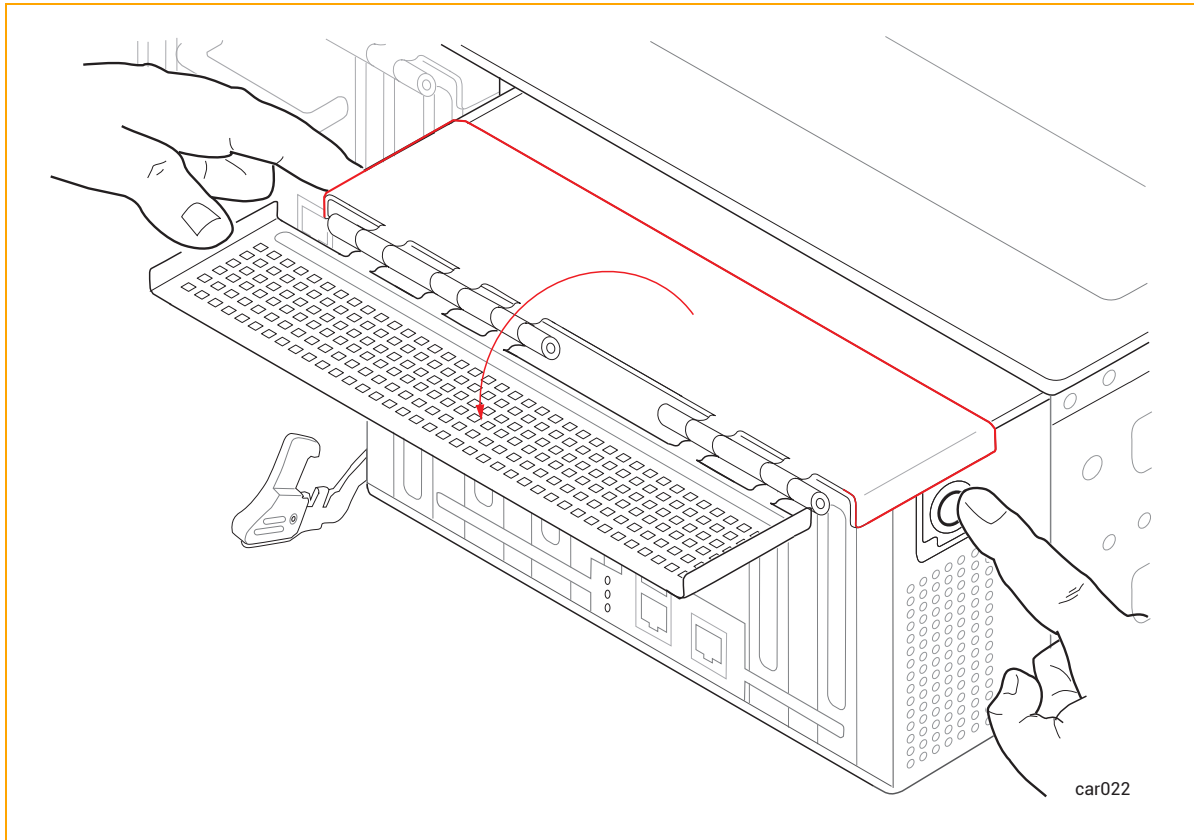
3. If your PCIe adapter contains SFP modules, remove them and go to the next step.

    Otherwise, go to the next step.

4. Pull the PCIe adapter out of its connector. See Figure 5-14.

Figure 5-14: Removing PCIe Adapter

car023
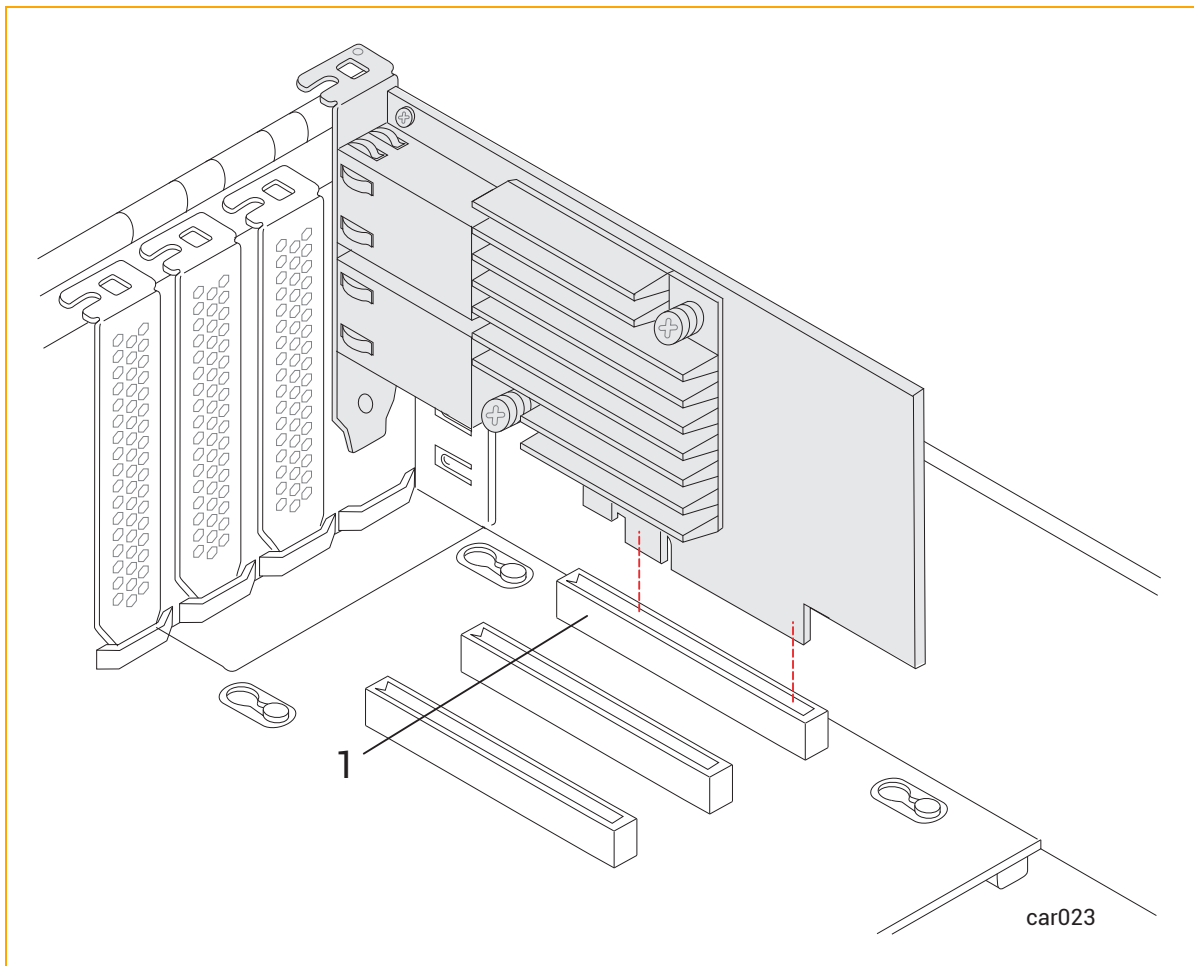
1	Connector

5. If the adapter does not disconnect easily, use a firm side-to-side tugging motion while maintaining steady pressure on the adapter.

To install a PCIe adapter

> **Note**: You should be familiar with the important information provided in "PCIe Adapters" on page 30 before you install a PCIe adapter.

> **Notice**: Before you begin this procedure, move the I/O module so that its front overhangs the surface on which it is resting by approximately two inches. Otherwise, the PCIe adapter cannot be installed correctly.

1. If the PCIe adapter contains SFP modules, remove them and go to the next step.

   Otherwise, go to the next step.

2. Install the PCIe adapter firmly into its connector. To do so, apply pressure to the part of the PCIe adapter located above the far side of the board connector (see arrow in Figure 5-15). Once installed, make sure that the PCIe adapter's top bracket is flush with the chassis.

   Figure 5-15: Installing PCIe Adapter

car024
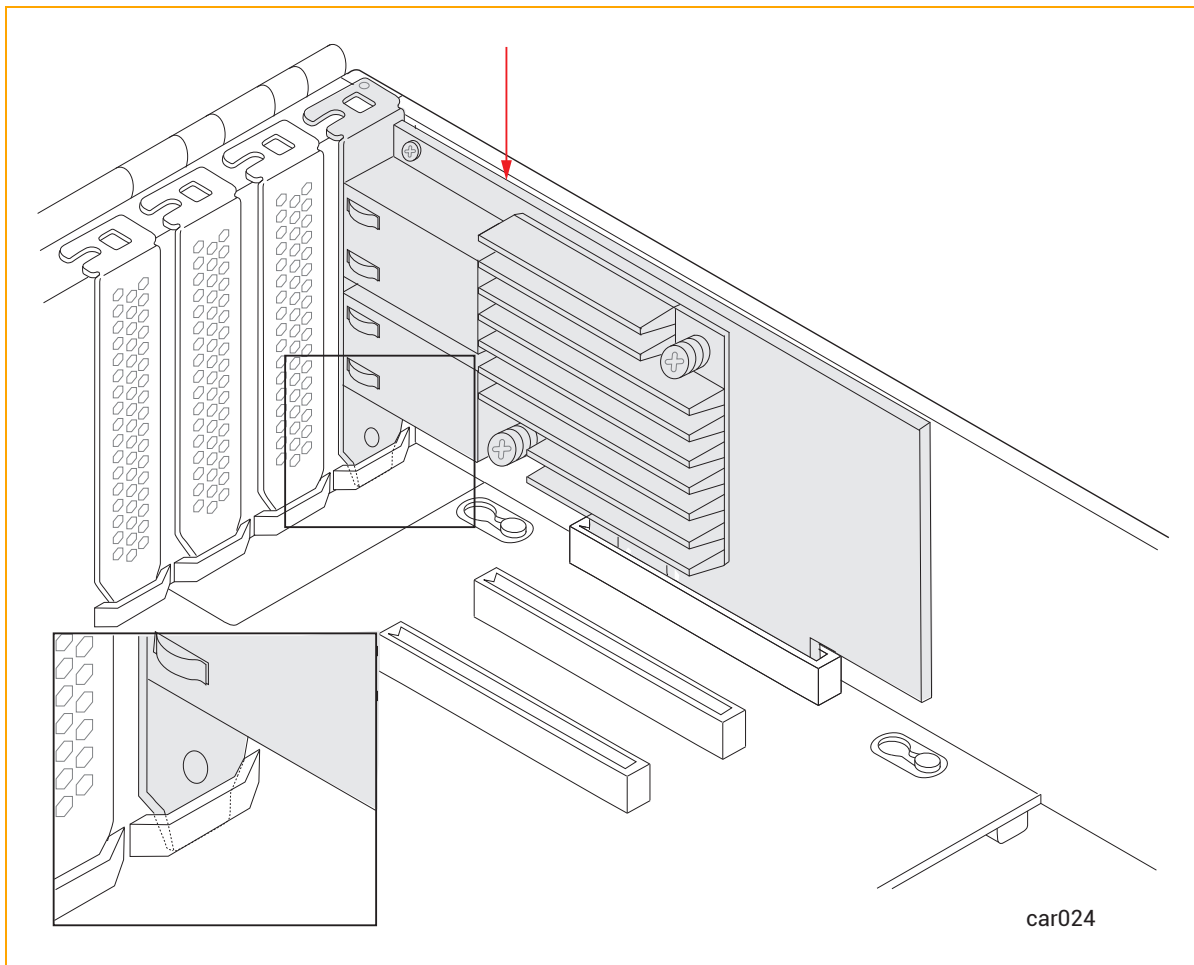
> **Notice**: If the PCIe adapter is not seated correctly, the metal cage cannot close properly. This may lead to the I/O module becoming lodged in the chassis.

3. If you removed SFP modules earlier, reinsert them in the PCIe adapter and go to the next step.

   Otherwise, go to the next step.

4. Flip the metal cage closed until it clicks. See Figure 5-16.

Figure 5-16: Closing Metal Cage

car028
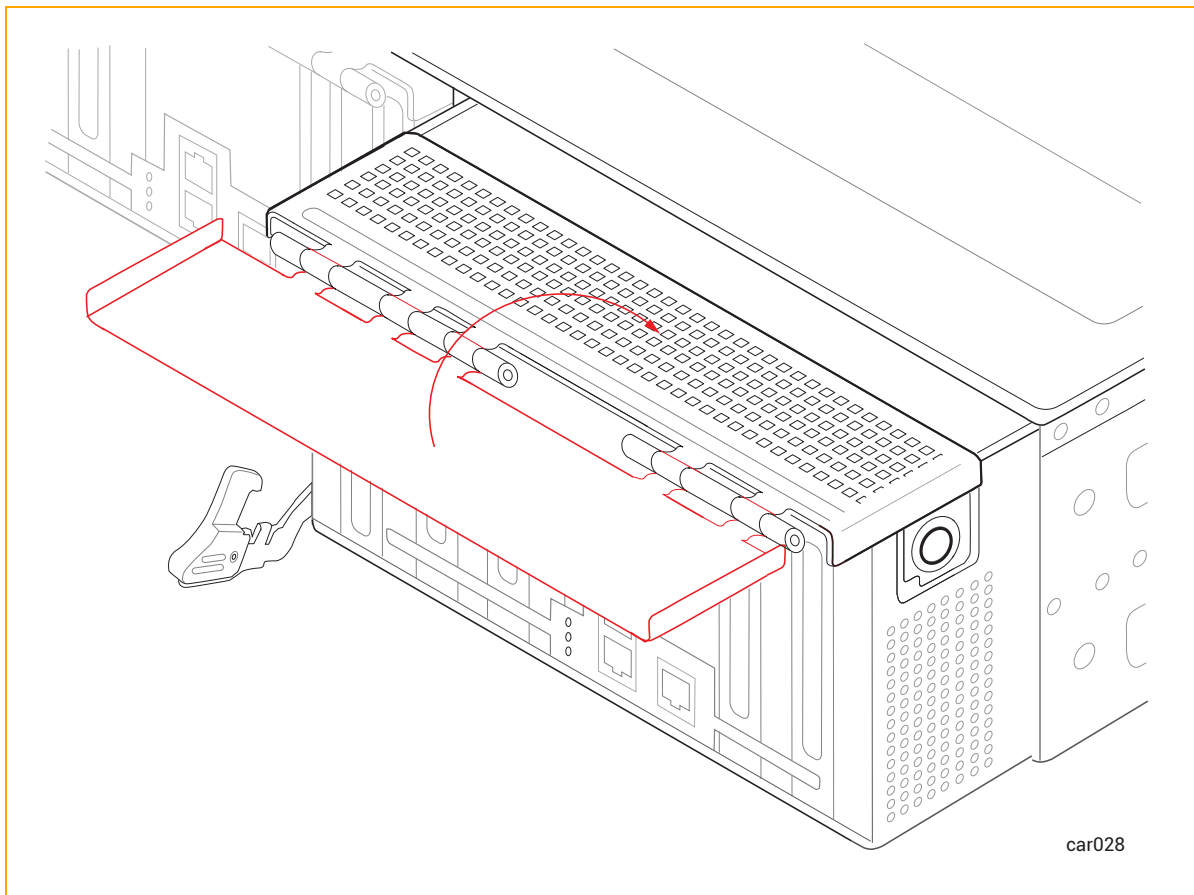
5. Replace the I/O module.

Related Topics

- "PCIe Adapters" on page 30

- "I/O Module" on page 29

Removing and Replacing a PSU and Power Cords

> **Cautions**:
>
> 1. Use ESD precautions when performing these procedures.
>
> 2. Read "Replacing System Components" on page 104 before you remove any system components.
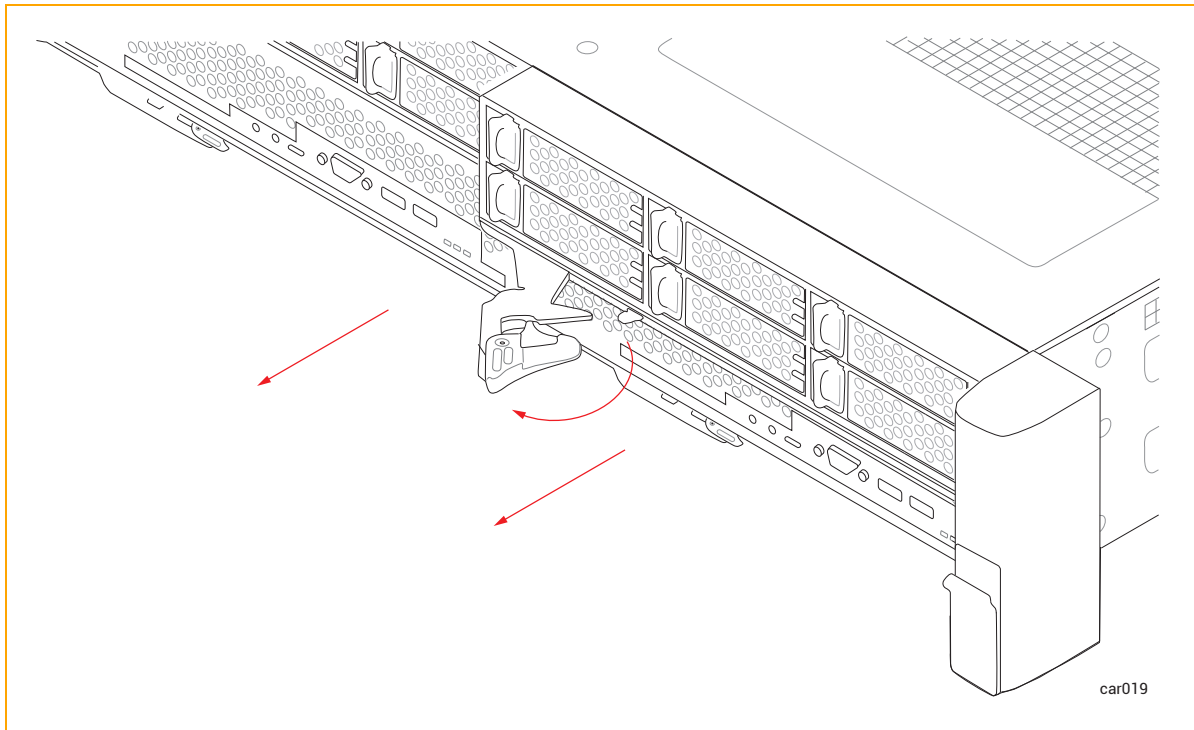
Précautions:

1. Utilisez les précautions ESD lors de l'exécution de ces procédures.

2. Lisez "Replacing System Components" on page 104 avant de retirer des composants du système.

To remove a PSU

1. Unwrap the power-cord retention strap, and then disconnect the power cord. See Figure 5-17.

Figure 5-17: Disconnecting Power Cord and Retention Strap

car039
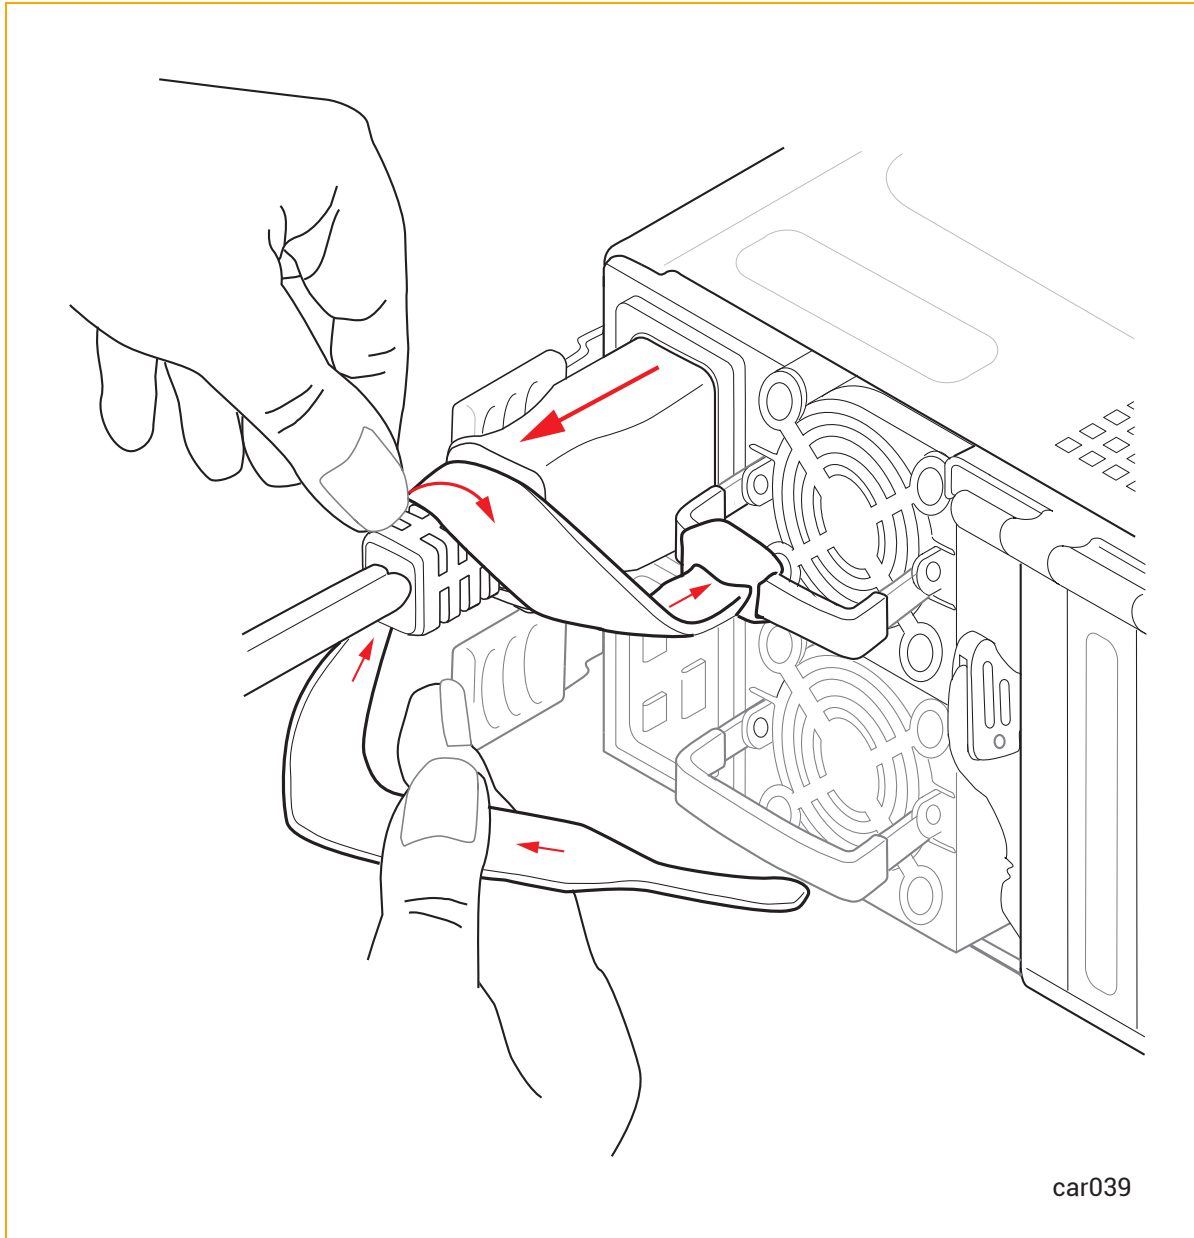
2. Push the release lever to the right while pulling the handle until the PSU releases. See Figure 5-18.

Figure 5-18: Pushing PSU Release Lever

car025
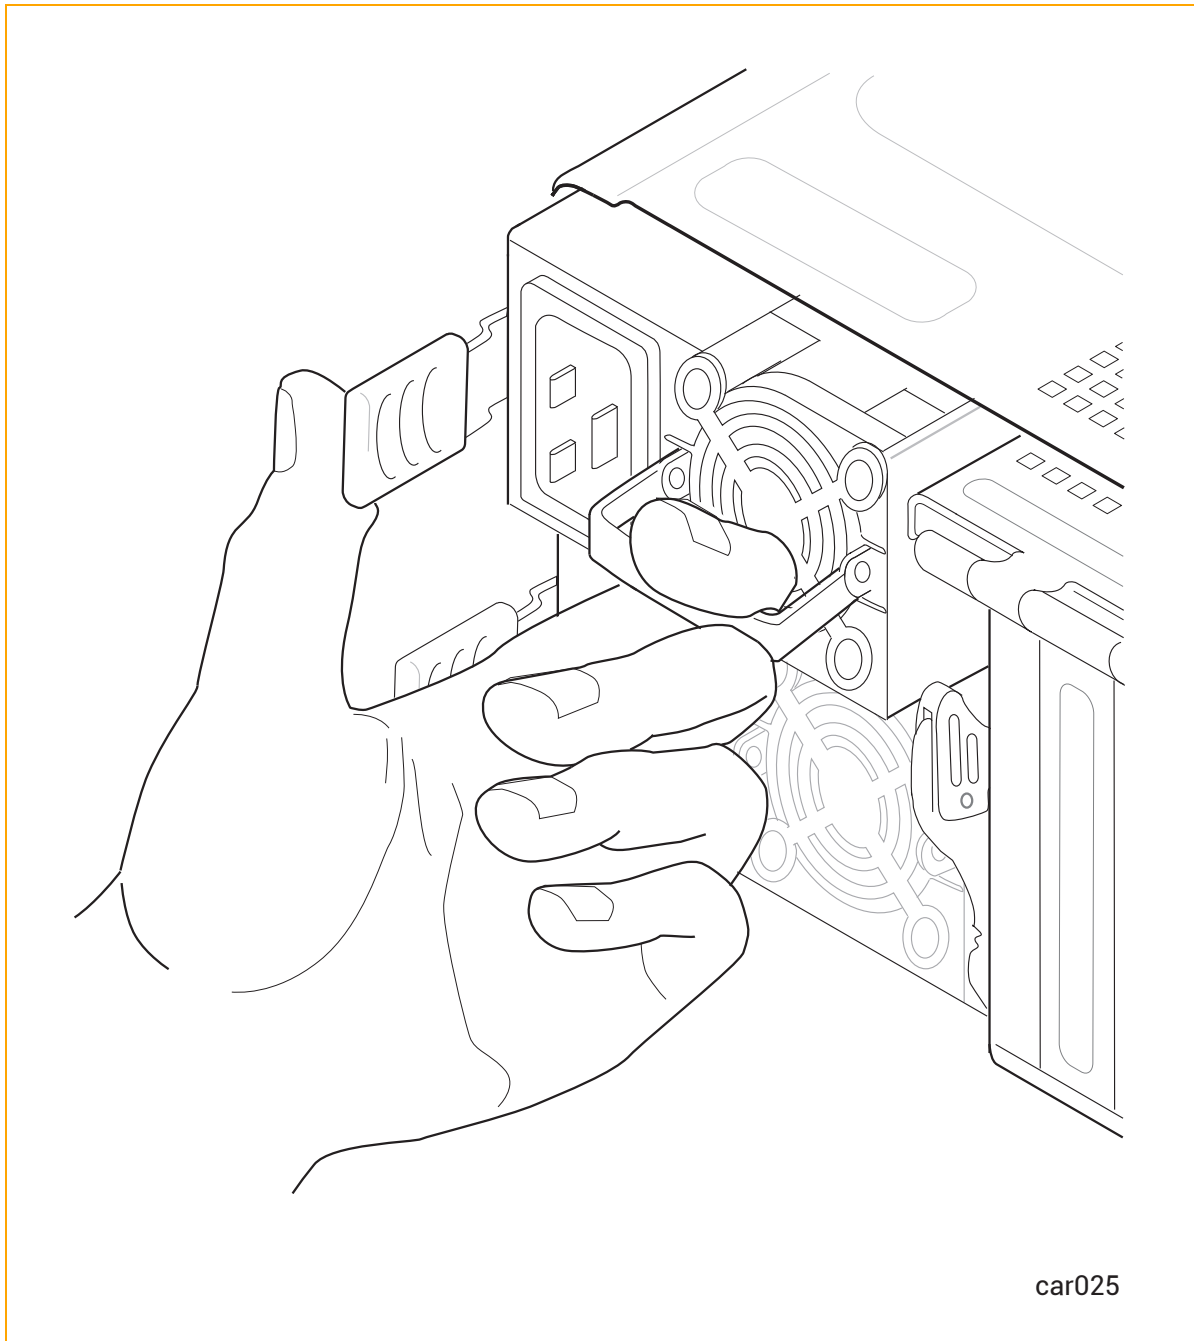
3. Release the release lever and pull the PSU out of the system chassis.

To replace a PSU

> **Caution**: When you insert a PSU into the system chassis, make sure that the plug is on the left and the fan is on the right, as shown in Figure 5-17. Otherwise, damage may result.

> **Précaution**: Lorsque vous insérez un bloc d'alimentation dans le châssis du système, assurez-vous que la fiche est à gauche et que le ventilateur est à droite, comme indiqué sur Figure 5-17. Sinon, des dommages pourraient en résulter.

1. Slide the PSU into the system chassis until it stops.

2. If necessary, push the release lever to the right and gently push the PSU in until it is fully seated and the release lever latches.

3. Connect the power cord, and perform the steps in "To fasten the power-cord retention strap" on page 78.

To replace both power cords

1. Power off the system as described in "Shutting Down the System" on page 87.

2. Disconnect all communications cables from the system.

3. Unwrap the power-cord retention straps, and then disconnect the power cords from the PSUs and the power sources.

4. Wait for all LEDs in the front and rear of the system to turn off, including the LEDs on both PSUs. All LEDs turning off indicates that the standby voltages have dissipated.

5. Replace the two power cords, using the correct power cords for the voltage being used, and fasten the power-cord retention straps.

6. Reconnect all communications cables to the system.

6

Chapter 6: Electrical Circuit and Wiring Information

For electrical circuit and wiring information that you need to provide to the contractor and/or facilities personnel responsible for wiring the power at the system installation site, see:

- "Fault Protection Requirements" on page 131
- "Grounding Considerations" on page 131
- "Circuit Wiring Diagrams" on page 132
- "Electrical Power Connectors" on page 138

Fault Protection Requirements

Each CRU module in a ztC Endurance system contains internal fault/overload current protection. However, the system relies on the power distribution system at your site for protection against potential faults in the power cords and the wiring in the PSU AC inlet.

Power cords for a ztC Endurance system and PDU must not be connected to a branch circuit greater than 20A.

Related Topics

- "Electrical Circuit and Wiring Information" on page 131

Grounding Considerations

The system obtains an earth reference ground through the power cords attached to the system. Similarly, each peripheral device connected to the system obtains ground through its power cord. You **must** ensure that a high-integrity safety-ground conductor is installed as part of the wiring system (in accordance with U.S. national electric code NFPA 70 or the equivalent). The international safety standard (EN62368) for audio/video, information, and communication technology equipment also requires a ground conductor, but calls it a protective earth (PE) ground.

Depending on local conditions, ground potentials may differ between the system base and any peripheral devices connected to the system base. All grounds in the system **must** return to the same reference point in the power distribution system, as close as possible to **zero (0) volt potential** relative to earth reference ground. Earth reference ground is typically a metal stake in the ground to which the ground conductors from one or more buildings are attached.

As shown in Figure 6-1, a *star ground* is often used to obtain the same earth reference ground. Each earth reference ground, such as the system base ground, is returned separately to a common point where a zero-volt (0V) earth ground exists. The star ground ensures that all equipment is at the same potential and that no noise or safety problems associated with an unpredictable or uncharacterized grounding system will occur.

Figure 6-1: Star Ground Example
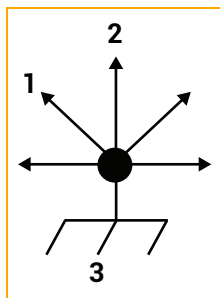
1	To monitor
2	To ztC Endurance system
3	Earth reference ground (0V)

Related Topics

- "Electrical Circuit and Wiring Information" on page 131

Circuit Wiring Diagrams

The circuit wiring diagrams shown in this topic illustrate how the hot, ground, and/or neutral AC signals should be connected to the power input plugs of the ztC Endurance system and other components.

In the following diagrams, the power inputs for ztC Endurance systems are labeled X and Y, as shown in Figure 6-2 and Figure 6-3, to eliminate any ambiguities in the nomenclature. For single-phase applications, the X input is connected to the L (Line) hot input, and the Y input is connected to the N (Neutral) input. However, for split-phase or three-phase applications, the X and Y inputs are connected to L1, L2, or L3 (separate lines). Therefore, for split-phase or three-phase applications, both X and Y can be electrically hot with respect to the system base (earth reference ground).

Figure 6-2 shows a face view of the physical locations of the X and Y inputs on a customer-supplied PDU (either A or B).

Figure 6-2: PDU Power Input Labeling
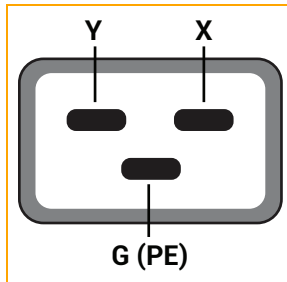
Figure 6-3 shows a face view of the physical locations of the X and Y inputs on a ztC Endurance system's PSU (either A or B).

Figure 6-3: ztC Endurance Enclosure Power Input Labeling
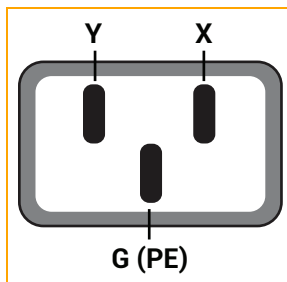
Figure 6-4 shows a single-phase 120V AC circuit connection. Note that this application requires a single-pole circuit breaker.

Figure 6-4: Single-Phase 120V AC Circuit Connection
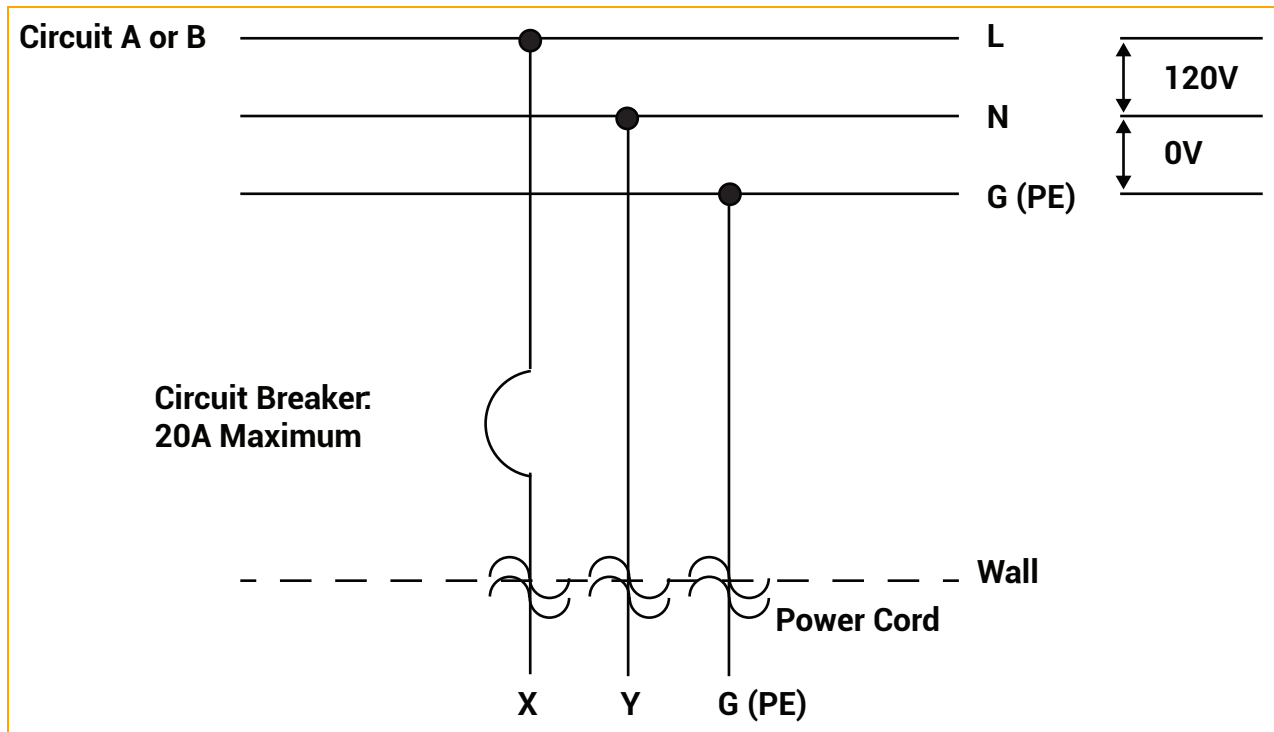
Figure 6-5 shows a single-phase 240V AC circuit connection. Note that this application requires a single-pole circuit breaker.

Figure 6-5: Single-Phase 240V AC Circuit Connection
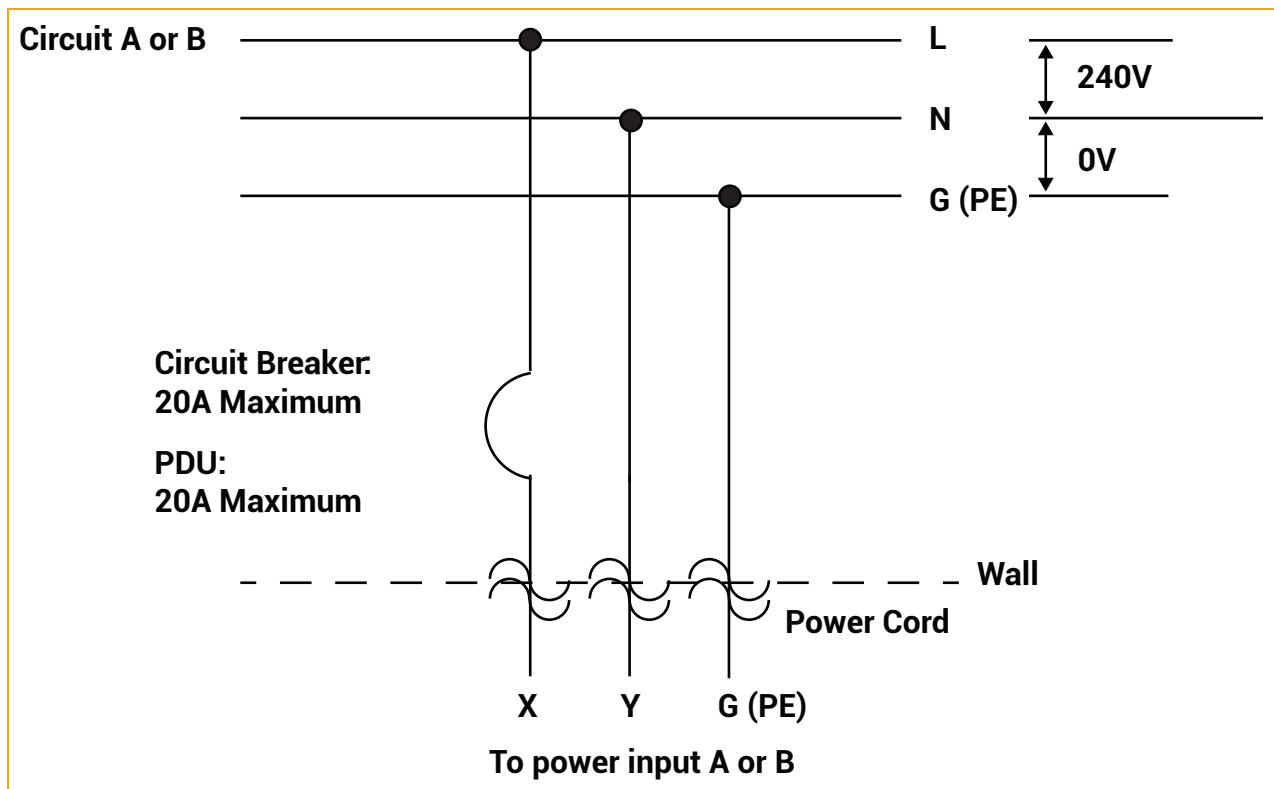
Figure 6-6 shows a split-phase 120/240V AC circuit connection. Note that this application requires a double-pole circuit breaker.

Figure 6-6: Split-Phase 120/240 Volts AC Circuit Connection
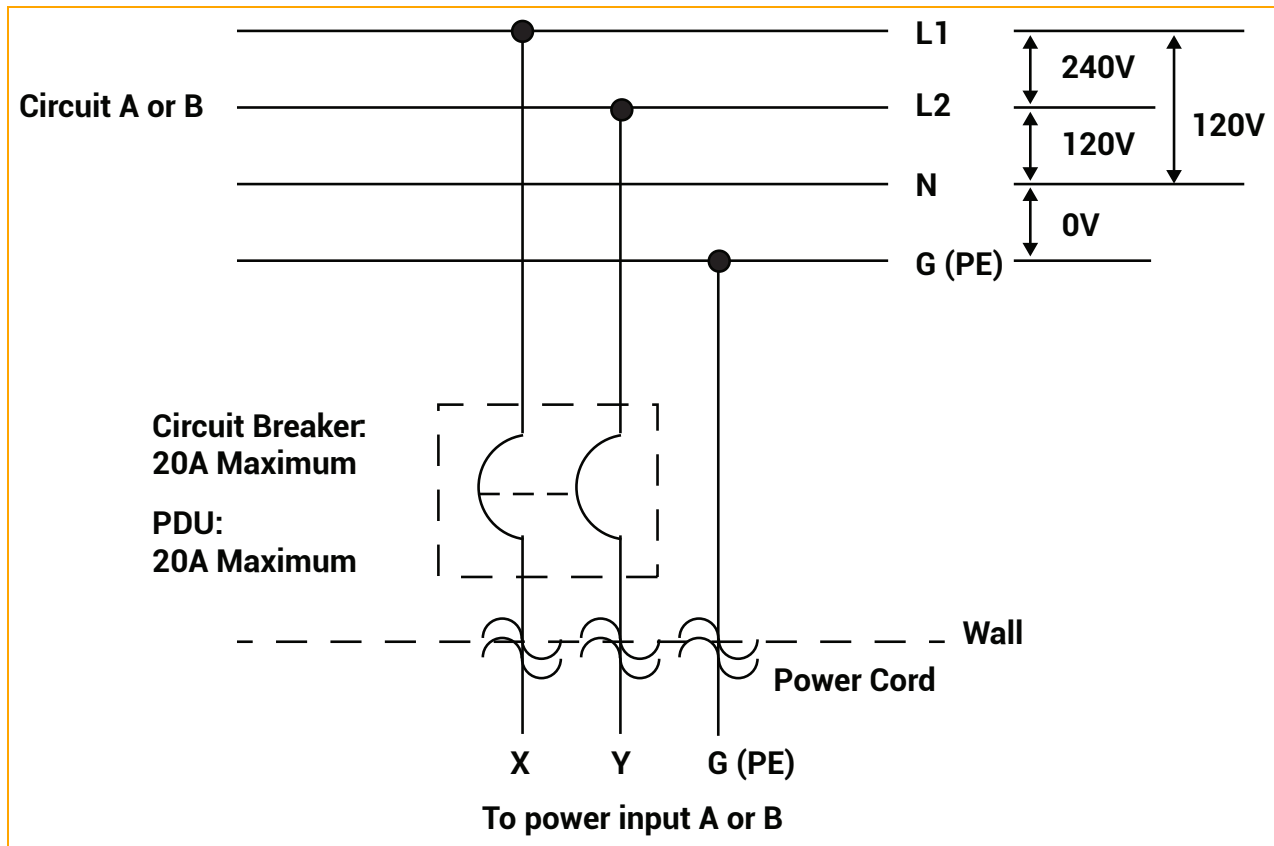
Figure 6-7 shows a three-phase 208V AC, Y-, or Δ-source circuit connection, which is a phase-to-phase source connection. Note that the X and Y inputs on the ztC Endurance system can be connected from L1 and L2, L2 and L3, or L1 and L3. This application requires a double-pole circuit breaker.

Figure 6-7: Three-Phase 208V AC, Y-, or Δ-Source Circuit Connection, Phase-to-Phase
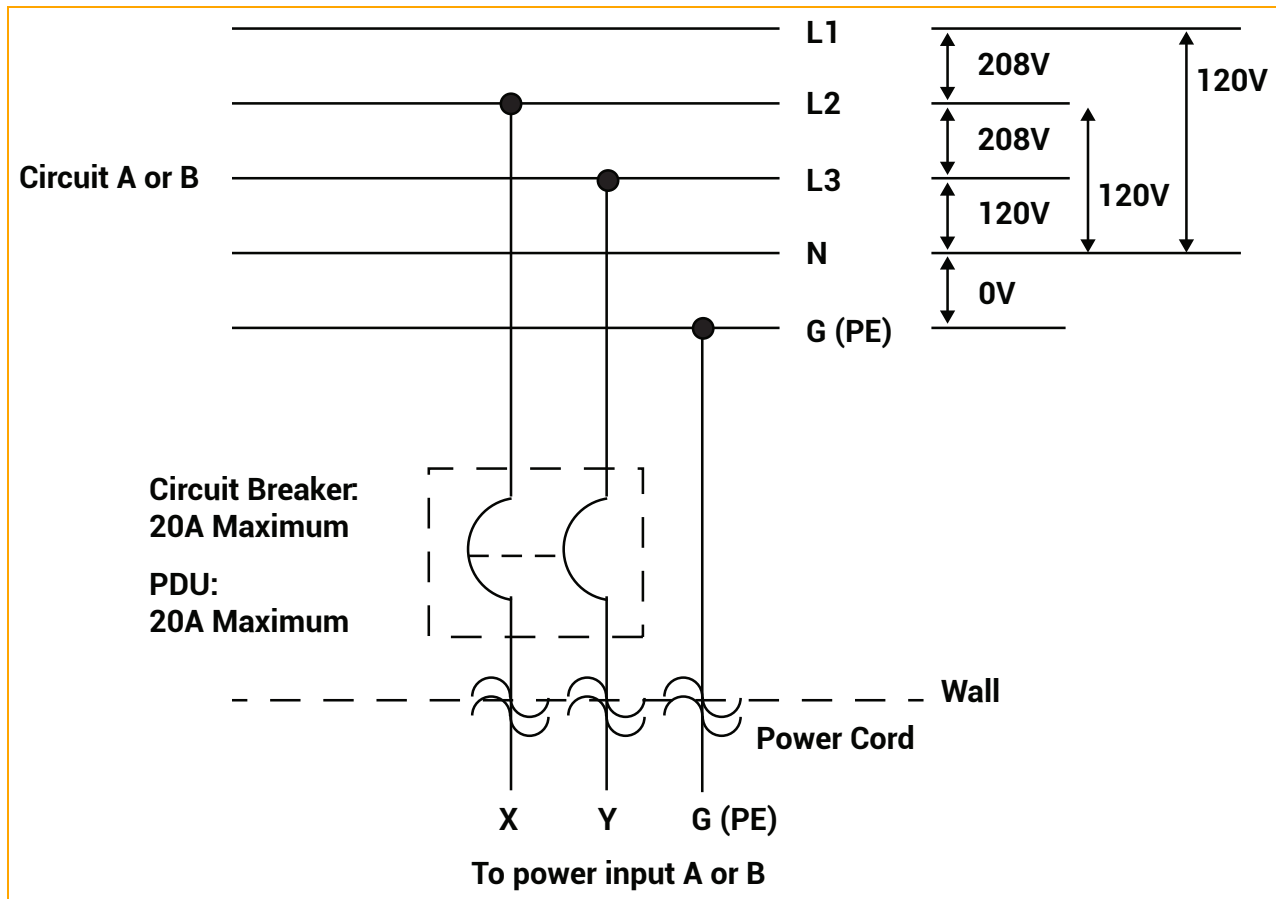
Figure 6-8 shows a three-phase 380V AC, Y-, or Δ-source circuit connection, which is a phase-to-neutral source connection. Note that the ztC Endurance system's X input can be connected to L1, L2, or L3. This application requires a single-pole circuit breaker.

Figure 6-8: Three-Phase 380V AC, Y-, or Δ-Source Circuit Connection, Phase-to-Neutral
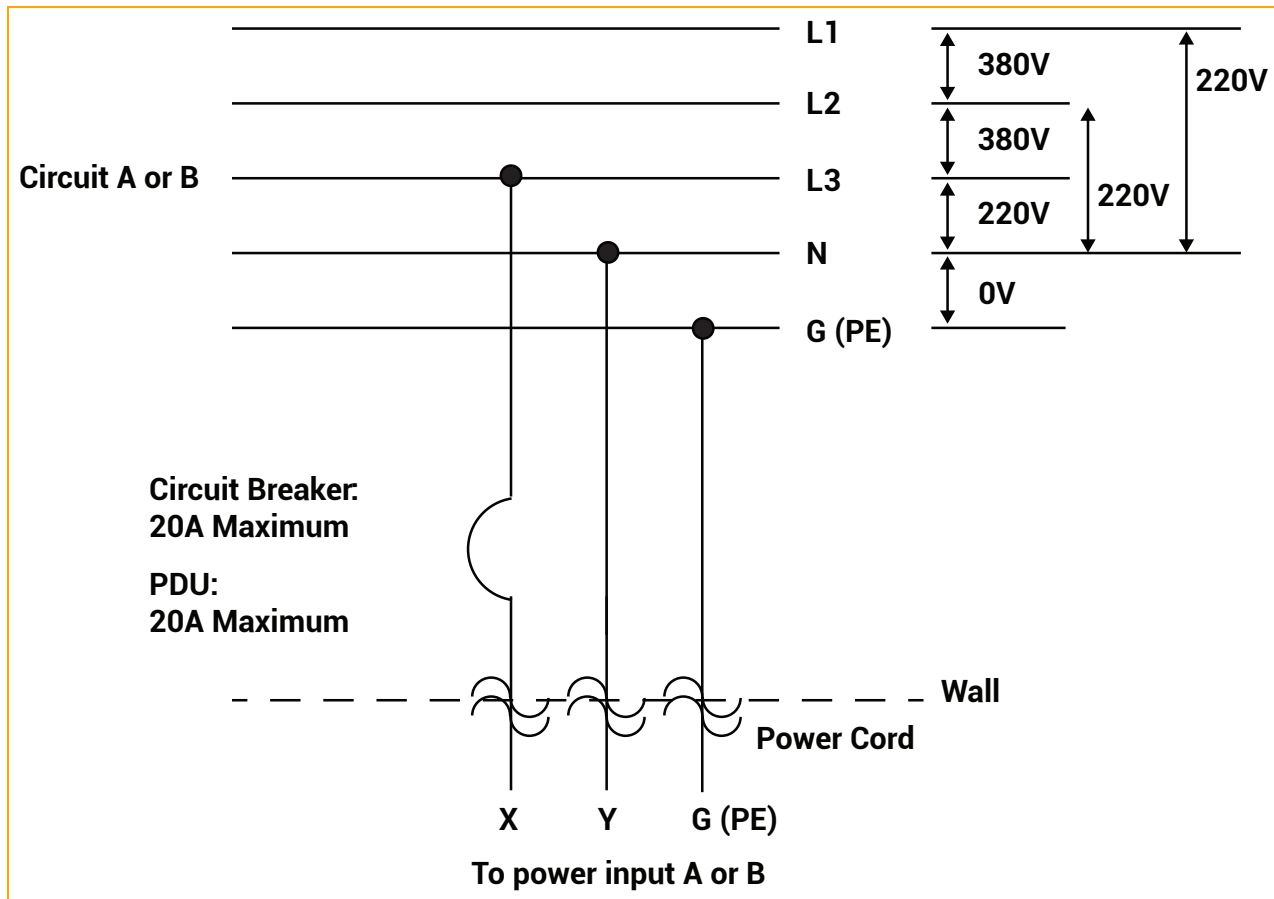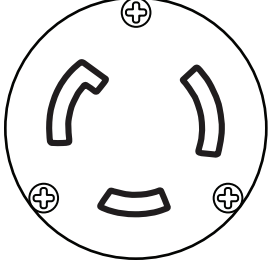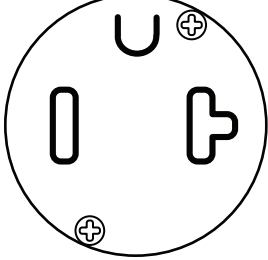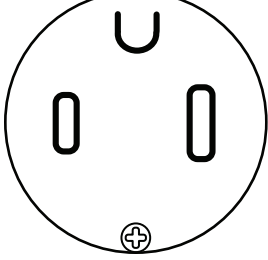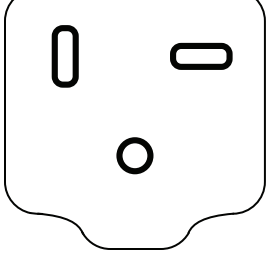
Related Topics
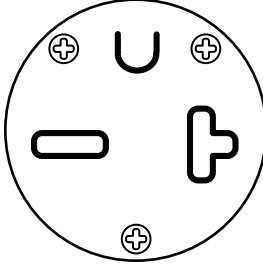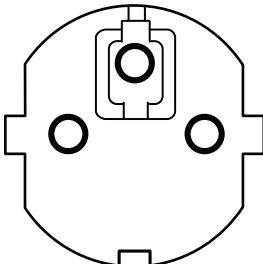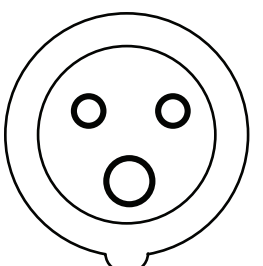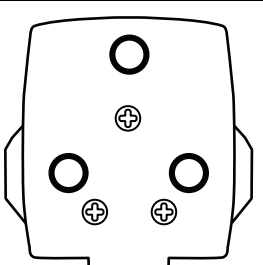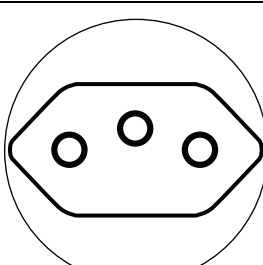
- "Electrical Circuit and Wiring Information" on page 131

Electrical Power Connectors
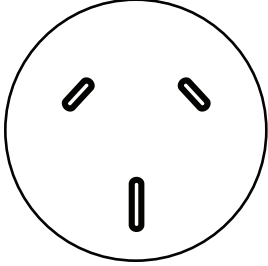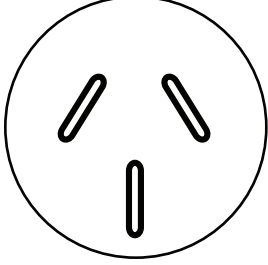
Table 6-1 describes the connectors (that is, the receptacles) required by the AC power cords that Stratus supplies with ztC Endurance systems and optional devices.

Table 6-1: Connectors for AC Power Outlets

Connector	Configuration	Rating	Description
NEMA L6-20		20A, 250 volts AC	2-pole, 3-wire
NEMA 5-20		20A, 125 volts AC	2-pole, 3-wire
JIS C 8303		15A, 125 volts AC	2-pole, 3-wire
JIS C 8303		15A, 250 volts AC	2-pole, 3-wire

Connector	Configuration	Rating	Description
NEMA 6-20 / CNS 690		20A, 250 volts AC	2-pole, 3-wire
CEE (7) VII		16A, 250 volts AC	2-pole, 3-wire
IEC 60309 (formerly IEC 309)		16-20A, 250 volts AC	2-pole, 3-wire
SABS 164-1:1992 / IS 1293 IND 16		16A, 250 volts AC	2-pole, 3-wire
NBR 14136 (Brazil)		10A, 250 volts AC	2-pole, 3-wire

Connector	Configuration	Rating	Description
SAA/3/15 AS/NZS 3112-1993		15A, 250 volts AC	2-pole, 3-wire
GB1002-1996		10A, 250 volts AC	2-pole, 3-wire

Related Topics

- "Electrical Circuit and Wiring Information" on page 131

Part 2: ztC Endurance Software Guide

The ztC Endurance Software Guide discusses the following topics:

Related Topics

Chapter 6: System Software Installation and Upgrade (Linux)

To install the system software, see "Installing the System Software" on page 143.

To upgrade the system software, see "Upgrading the System Software" on page 188.

Installing the System Software

> **Notice**: Before you install the system software, always check the Release Notes: Automated Uptime Layer with Smart Exchange Release RH-2.0.0.0 (Linux) for any installation-related issues.

Before you install the system software, you must perform some tasks to prepare for the installation. One of the tasks you perform is to determine if Stratus preinstalled the system software (that is, Automated Uptime Layer with Smart Exchange and Red Hat Enterprise Linux, or RHEL) at the factory. This information determines which type of system software installation you perform:

- If Stratus preinstalled the system software at the factory, you perform a *factory-installed system software installation*.

- If Stratus did not preinstall the system software at the factory, you perform a *full system software installation*. This installation procedure installs Automated Uptime Layer with Smart Exchange and RHEL, and allows you to enter customized configuration settings for your ztC Endurance system.

> **Note**: On all ztC Endurance systems, the Standby OS (Ubuntu, a Linux distribution) is preinstalled at the factory in an internal NVMe drive inside each compute module. In general, you need to access the Standby OS only for a full system software installation or troubleshooting purposes.

For both types of installations, you power on the system to boot the preinstalled Standby OS on both compute modules, and you later run configuration scripts to enter different settings that are required to support the software installation.

After the software installation completes, you need to perform some post-installation tasks, as described in "Post-Installation Tasks" on page 183.

Next Steps

Prepare for the system software installation as described in "Preparing to Install the System Software" on page 144.

Preparing to Install the System Software

Perform the following steps before you install the system software:

"I. Review the Preinstallation Checklist" on page 145

"II. Determine Which Software Installation Procedure to Perform" on page 157

I. Review the Preinstallation Checklist

Review the following checklist before you install the system software.

> **Caution**: Make sure that the network switch used during the system software installation is 1 Gb or faster. Otherwise, the copy of the system software ISO to the host system may time out, resulting in a failed installation.

Preinstallation Checklist

☐ Gather and record the site-specific settings described in "Gathering Information" on page 146. The installation and configuration scripts require this information.

Gathering Information

Review the information in Table 6-1 and Table 6-2, and then print the "Password and Installation Settings Printable Worksheet" on page 158. Use it to record your passwords and installation settings.

> **Notice**: Keep your passwords and installation settings in a secure place so that you can access your ztC Endurance system after installation. Keeping this information accessible also helps your authorized Penguin Solutions service representative to access the system in the event of a service call.

> **Note**: During the software installation, you do not set passwords for the BMC Debug console `sysadmin` or BMC Redfish `Administrator` user accounts. If you need to change these passwords at a later time, use the **Change Passwords** page on the ztC Endurance console. See "Changing Passwords on a ztC Endurance System" on page 228 for more information.

Table 6-1 describes the user-configurable passwords required for the software installation. You are prompted for these values only on the command line.

Table 6-1: Password Settings

Setting	Description
System Passwords	
BMC Web console admin password	Password for the BMC Web console `admin` user account.
	You need to set a password for the BMC Web console to support the software installation and troubleshooting, but you can also use the BMC Web console to remotely monitor and administer the system as summarized in "Configuring the BMCs for Remote System Management" on page 186.
	The password, which is case-sensitive, must be a string of 8 to 16 alphanumeric characters (with an exception for the default password of `admin`). Spaces are not allowed. The

Setting	Description
	following special characters are allowed: - (hyphen), _ (underscore), @ (at-sign), % (percent sign), & (ampersand), \ (backslash), . (period),* (asterisk).
Management zenadmin password	Password for the `zenadmin` user account of the ztC Endurance Standby operating system (OS).
	You log in to the Standby OS to install the system software, as described in "Performing a Full System Software Installation" on page 166. Specifying a new password for the `zenadmin` account during the software installation replaces the default password (`zenadmin`).
	The password, which is case-sensitive, must be a string of 6 to 20 alphanumeric characters and must contain both upper- and lowercase alphabetical characters. It cannot be the same as the username. The following special characters are allowed: - (hyphen), _ (underscore), @ (at-sign), % (percent sign), & (ampersand), # (pound sign), . (period), $ (dollar sign), * (asterisk), ! (exclamation point).
RHEL admin-istrator password	Password for the Administrator account of the RHEL host.
	The password, which is case-sensitive, must be a string of 8 to 24 alphanumeric characters and must contain both upper- and lowercase alphabetical characters and a numeric character. It cannot be the same as the username. The following special characters are allowed: - (hyphen), _ (underscore), @ (at-sign), % (percent sign), & (ampersand), # (pound sign), \ (backslash), . (period), $ (dollar sign), * (asterisk).

Table 6-2 describes the system and network settings needed for the software installation. The type of system software installation you are performing determines how you supply the settings:

- **Command-line prompts**. If Stratus installed the system software at the factory, you must supply the settings on the command line. If you are performing a full system software installation, you can optionally supply the settings on the command line.

- **Configuration-file entries**. If Stratus did not install the system software at the factory, you can optionally supply the settings in an installation configuration file (`zen_config.yaml`). Using a installation configuration file saves you time, as you can reuse the file to run other configuration scripts. If Stratus installed the system at the factory, you cannot use the installation configuration file to supply the settings.

Most settings in Table 6-2 show two **Setting** names:

- Longer setting names in mixed case (for example, **BMCA hostname**) are similar to the command-line prompts in the installation script.

- Shorter setting names in all lowercase letters (for example, **bmca_ hostname**) match the names in the installation configuration file. Table 6-2 may not reflect the exact order in which you are asked to supply the settings in the file.

> **Note**:
> A ztC Endurance system requires a total of 11 IP addresses for the following interfaces in order to provide and configure full redundancy:
>
> - Ten interfaces on the management network:
>
>   - One for each BMC port (four total)
>
>   - One for each embedded **1G MGMT** port for use by each Standby OS (four total)
>
>   - One for each embedded **1G MGMT** port for use by the RHEL OS (two total)
>
> - One interface on the production network, for the bond0 interface of embedded **10G PORT2** ports in each I/O module
>
> You can configure each interface type to use either the static or DHCP protocol. Also, you can configure additional interfaces after the installation process is complete.

Table 6-2: Installation Settings

Setting	Description
System Information Variables	
system_id	**Note: Applies only if you are preparing an installation configuration file.** The system ID to include in host name construction. For example, if the fully qualified domain name (FQDN) of a system is `ocean.abc.com`, where `ocean` is the host name and the DNS domain name is `abc.com`, the system_id is `ocean`. This setting is helpful if you are creating installation configuration files for many systems, so you can enter the ID only once at the top of the configuration file, and replace it with the `$SYSTEM_ID` variable wherever necessary.
system_domain	**Note: Applies only if you are preparing an installation configuration file.** The system domain to include in host name construction. For example, if the fully qualified domain name of a system is `ocean.abc.com`, where `ocean` is the host name and the DNS domain name is `abc.com`, the system domain is `abc.com`. This setting is helpful if you are creating installation configuration files for many systems, so you can enter the domain only once at the top of the configuration file, and replace it with the `$SYSTEM-DOMAIN` variable wherever necessary.
system_ nameserver	**Note: Applies only if you are preparing an installation configuration file.** The system name server to use in name server fields in the installation configuration file. Specify an IP address for this variable, or specify `--NONE--`.

Setting	Description
	This setting is helpful if you are creating installation configuration files for many systems, so you can enter the name server only once at the top of the configuration file, and replace it with the $NAMESERVER variable wherever necessary.
BMC Settings	
BMC A hostname bmca_hostname	BMC A host name, where BMC A is the BMC in compute module A. You can specify either the host name or the fully qualified domain name) (FQDN). For example: `sys20bmca`. You connect to this host name or the IP address of BMC A to remotely monitor and manage compute module A with the BMC Web console. For more information, see "Configuring the BMCs for Remote System Management" on page 186.
BMC B hostname bmcb_hostname	BMC B host name, where BMC B is the BMC in compute module B. You can specify either the host name or the FQDN. For example: `sys20bmcb`. You connect to this host name or the IP address of BMC B to remotely monitor and manage compute module B with the BMC Web console.For more information, see "Configuring the BMCs for Remote System Management" on page 186.
BMC Network protocol bmc_netproto	BMC IP protocol (`dhcp` or `static`).
BMC A eth0 IP bmca_eth0	BMC A eth0 IPv4 or IPv6 address (IP address/prefix) or leave blank (for DHCP).

Setting	Description
BMC A eth1 IP bmca_eth1	BMC A eth1 IPv4 or IPv6 address (IP address/prefix) or leave blank (for DHCP).
BMC B eth0 IP bmcb_eth0	BMC B eth0 IPv4 or IPv6 address (IP address/prefix) or leave blank (for DHCP).
BMC B eth1 IP bmcb_eth1	BMC B eth1 IPv4 or IPv6 address (IP address/prefix) or leave blank (for DHCP).
Standby OS Settings	
Standby A Hostname standbya_ hostname	Standby A host name, where Standby A is the Standby OS running on compute module A. You can specify either the host name or the FQDN. For example: `sys20a`.
Standby B Hostname standbyb_ hostname	Standby B host name, where Standby B is the Standby OS running on compute module B. You can specify either the host name or the FQDN. For example: `sys20b`.
Standby Network protocol standby_netproto	Standby IP protocol (`dhcp` or `static`).
Standby A IP Address 1 standbya_ ipaddr_1	Standby A IPv4 or IPv6 address (IP address/prefix) or leave blank (for DHCP) for A eno1.
Standby A IP Address 2	Standby A IPv4 or IPv6 address (IP address/prefix) or leave blank (for DHCP) for A eno2.

Setting	Description
standbya_ ipaddr_2	
Standby B IP Address 1 standbyb_ ipaddr_1	Standby B IPv4 or IPv6 address (IP address/prefix) or leave blank (for DHCP) for B eno1.
Standby B IP Address 2 standbyb_ ipaddr_2	Standby B IPv4 or IPv6 address (IP address/prefix) or leave blank (for DHCP) for B eno2.
Standby Nameserver standby_ nameserver	Standby IPv4/IPv6 name server address or leave blank (for DHCP). **Note for installation configuration file users:** If you do not want to use a name server and do not want to be prompted for one during the installation, specify `--NONE--`.
RHEL Settings	
RHEL Hostname rhel_hostname	Name of the RHEL host. You can specify either the host name or the FQDN. For example: `sys20`. You connect to this host name or the IP address of your RHEL host to configure and manage your system.
RHEL Network Protocol rhel_bootproto	RHEL host IP protocol (`dhcp` or `static`).
RHEL IP Address rhel_ipaddr	RHEL host IP address.

Setting	Description
RHEL Gateway rhel_gateway	RHEL host IPv4/IPv6 gateway address or leave blank (for DHCP).
RHEL Netmask rhel_netmask	RHEL host IPv4 subnet address or leave blank (for DHCP/IPv6 or if prefix is supplied in IP address).
RHEL IPv6 Prefix rhel_ipv6_prefix	RHEL host IPv6 prefix or leave blank (for DHCP). Value must be an integer less than or equal to 128.
RHEL Nameserver rhel_nameserver	RHEL host IPv4/IPv6 name server address or leave blank. **Note for installation configuration file users:** If you do not want to use a name server and do not want to be prompted for one during the installation, specify `--NONE--`.
Path/URL to the RHEL ISO rhel_iso	**Note: Applies only if you are performing a full software installation.** RHEL ISO path or URL. For example, if you copy the ISO into the Standby OS as described later in "To transfer the RHEL ISO image to the ztC Endurance system with a USB storage device" on page 173, the file path would be similar to the following: `/opt/stratus/iso/rhel-9.4-x86_64-dvd.iso`
RHEL Management Network Settings (Used by ztC Endurance console)	
RHEL Management Network Protocol rhel_mgmt_netproto	RHEL management IP protocol (`dhcp` or `static`).

Setting	Description
RHEL Management IP Address 1 rhel_mgmt_ ipaddr1	RHEL management network IPv4 subnet address or leave blank (for DHCP/IPv6 or if prefix is supplied in IP address).
RHEL Management IP Address 2 rhel_mgmt_ ipaddr2	RHEL management network IPv4 subnet address or leave blank (for DHCP/IPv6 or if prefix is supplied in IP address).
ztC Endurance Management Network Settings (Used by BMC and Standby OS)	
Zen Management Gateway Address zenmgmt_ gateway	Management network IPv4/IPv6 gateway address or leave blank (for DHCP).
Zen Management Netmask zenmgmt_ netmask	Management network IPv4 subnet address or leave blank (for DHCP/IPv6 or if prefix is supplied in IP address).
Mgmt IPv6 Prefix zenmgmt_ipv6_ prefix	Management network IPv6 prefix or leave blank (for DHCP/IPv6 or if prefix is supplied in IP address).
Support Notifications over the ActiveService Network (ASN) Settings	
Enable support notifications	Indicates whether or not to enable Support Notifications over the ASN (True or False).

Setting	Description
asn_enabled	Activating the Support Notifications allows the ztC Endurance system to send health and status notifications to Penguin Solutions Customer Service over a secure Internet connection. To access related settings in ztC Endurance console after installation, see "Configuring Remote Support Settings" on page 264.
Asset ID asn_asset_id	Asset Identifier (ID), in the format ze*nnnnn* (for example, zen12345). This value is required even if you do not enable Support Notifications. The asset ID is located in your sales order form and on top of the system chassis. For information about locating the asset ID sticker for your system on top of the system chassis, see "Installing the System in a Cabinet" on page 62.

☐ Make sure you have a **Penguin Solutions Customer Service Portal** account. If you do not already have an account, register for one at https://service.stratus.com.

☐ Check that you have the most current Release Notes: Automated Uptime Layer with Smart Exchange Release RH-2.0.0.0 (Linux) and software installation topics for your distribution. To restore your installation, make sure you have the appropriate version of the release level that you will restore.

☐ Prepare to accept the end-user license agreements (EULAs) that are presented during the installation process. To locate and read the text of the EULAs you must accept before installing Automated Uptime Layer with Smart Exchange, see "End-User License Agreements and Warranty" on page 300. You should not perform the installation if you cannot accept the EULAs or are not authorized to accept them. The installation process terminates if you decline any EULAs.

☐ Consider whether you want to enable Support Notifications over the ASN. Activating Support Notifications over the ASN allows the ztC Endurance system to send health and status notifications to Penguin Solutions Customer Service over a secure Internet connection.

You are prompted to enable or disable Support Notifications over the ASN during the software installation.

☐ With the ztC Endurance system halted, detach most peripheral devices from the system (exceptions are noted below). Devices to detach include:

- Unsupported PCIe adapters

- External storage systems

- USB devices, **except for** the keyboard, mouse, and USB storage device that you use to copy the RHEL ISO during a full system software installation

> **Notice**: The installation process cannot accommodate the presence of customer-added or unknown hardware. Add any such hardware, and configure the system to support it, only after completing the installation and verifying that the system functions as expected.

II. Determine Which Software Installation Procedure to Perform

The software installation procedure you perform depends on whether your company opted to have Stratus install the system software at the factory.

If you do not know whether Stratus installed the system software at the factory, check the sales order. If you do not have the sales order, perform the following procedure.

To determine which software installation procedure to perform

1. Turn on the monitor(s) and press the power button on the control-panel ear of the system (see Figure 6-1 for the location of the power button).

Figure 6-1: Power Button

car034
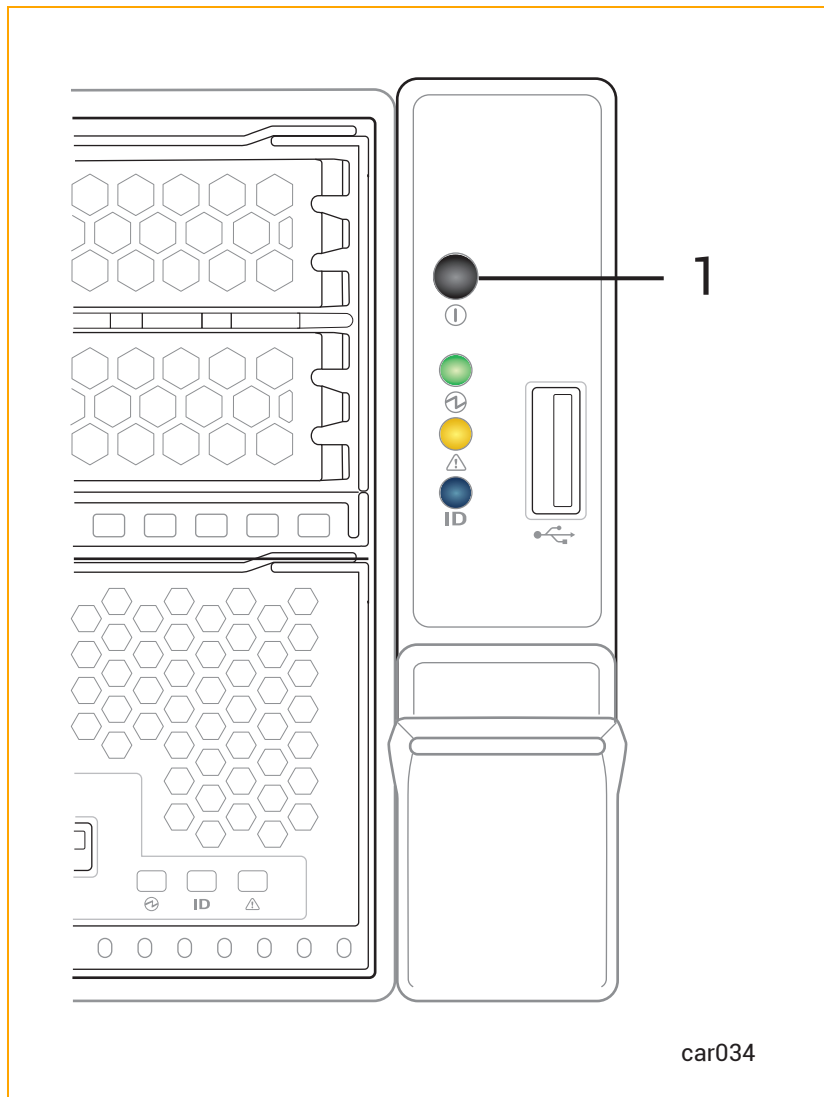
2. When the system boots, determine which of the following scenarios applies to your system:

   - **The system software was installed at the factory** if the system automatically boots the Standby OS on both modules, reboots the RHEL OS on compute module A, and the **Welcome to Red Hat Enterprise Linux** *release_number*! screen appears.

   - **The system software was not installed at the factory** if the system automatically boots the Standby OS on both modules and then does nothing else.

3. Power down the system by briefly pressing the power button again.

Next Steps

Perform the appropriate system software installation procedure:

- If Stratus installed the system software at the factory, see "Starting a Factory-Installed System for the First Time" on page 161.

- If Stratus did not install the system software at the factory, see "Performing a Full System Software Installation" on page 166.

Password and Installation Settings Printable Worksheet

Print the following worksheet and use it to record your passwords and installation settings. Be sure to save this checklist in a secure and accessible location.

After printing this worksheet, return to "Gathering Information" on page 146.

Table 6-3: Passwords and Installation Settings Worksheet

Setting	Entry
System Passwords	
BMC Web console admin password	
Management zenadmin password	
RHEL administrator password	
System Information Variables	

Setting	Entry
system_id	
system_domain	
system_nameserver	
BMC Settings	
BMC A hostname / bmca_hostname	
BMC B hostname / bmcb_hostname	
BMC Network protocol / bmc_netproto	
BMC A eth0 IP / bmca_eth0	
BMC A eth1 IP / bmca_eth1	
BMC B eth0 IP / bmcb_eth0	
BMC B eth1 IP / bmcb_eth1	
Standby OS Settings	
Standby A Hostname / standbya_hostname	
Standby B Hostname / standbyb_hostname	
Standby Network protocol / standby_netproto	
Standby A IP Address 1 / standbya_ipaddr_1	
Standby A IP Address 2 / standbya_ipaddr_2	

Setting	Entry
Standby B IP Address 1 / standbyb_ipaddr_1	
Standby B IP Address 2 / standbyb_ipaddr_2	
Standby Nameserver / standby_nameserver	
RHEL Settings	
RHEL Hostname / rhel_hostname	
RHEL Network Protocol / rhel_bootproto	
RHEL IP Address / rhel_ipaddr	
RHEL Gateway / rhel_gateway	
RHEL Netmask / rhel_netmask	
RHEL IPv6 Prefix / rhel_ipv6_prefix	
RHEL Nameserver / rhel_nameserver	
Path/URL to the RHEL ISO / rhel_iso	
RHEL Management Network Settings (Used by ztC Endurance console)	
RHEL Management Network Protocol / rhel_mgmt_netproto	
RHEL Management IP Address 1 / rhel_mgmt_ipaddr1	
RHEL Management IP Address 2 / rhel_mgmt_ipaddr2	

Setting	Entry
ztC Endurance Management Network Settings (Used by BMC and Standby OS)	
zen Management Gateway Address / zenmgmt_gateway	
zen Management Netmask / zenmgmt_netmask	
Mgmt IPv6 Prefix / zenmgmt_ipv6_prefix	
Support Notifications over the ActiveService Network (ASN) Settings	
Enable support notifications / asn_enabled	
Asset ID/ asn_asset_id	

Starting a Factory-Installed System for the First Time

If Stratus installed the system software at the factory, perform the following procedure to start the ztC Endurance system for the first time.

> **Notice**: Do not perform this procedure until you have performed all of the procedures described in "Preparing to Install the System Software" on page 144.

To complete the installation

1. Turn on the monitor(s) and press the power button on the control-panel ear of the system.

   > **Note**: If you are using only one monitor and keyboard, make sure they are connected to compute module A.

2. As each compute module boots, perform the following steps on **both** compute modules:

   a. When the Stratus ztC Endurance screen appears, press either the Delete key or the Esc key until the BIOS setup utility appears.

   b. Click **Server Mgmt**.

   c. Under **Server Mgmt**, click **BMC Network Configuration**.

d. Scroll to **LAN Channel 1** or **LAN Channel 8** and locate **Station IP Address**. Copy the IP address, paste it into your Web browser, and press **Enter** to connect to the BMC Web console.

e. Log in to the BMC Web console by typing the credentials (`admin` and `admin`, if this is your first time logging in, and click **Sign me in**. (If you need more information, see "Logging In to the BMC Web Console" on page 273.)

f. From the BMC Web console dashboard, click **Remote Control**, then click **Launch H5Viewer** (that is, the KVM).

3. From the KVM, monitor the boot process on the active compute module (**compute module A**). Compute module A automatically boots three times, and compute module B automatically boots once. Do not perform any actions until the **Welcome to Red Hat Enterprise Linux** *release_ number*! screen appears.

4. **On compute module A**:

a. Log in to the RHEL Administrator account, using the credentials you wrote down earlier on the "Password and Installation Settings Printable Worksheet" on page 158.

> **Notice**: The installation fails if you log in to any account other than the RHEL Administrator account during this step. If this occurs, contact Penguin Solutions Customer Service.

b. After the installation pauses for approximately one minute, type the `sudo /opt/stratus/bin/sra_sysconfig` command to start the Stratus ztC Endurance System Configuration Wizard.

> **Notice**: While the `sra_sysconfig` command is running, make sure a KVM is running for **both** compute modules. A Smart Exchange occurs approximately halfway through the process, and you can observe this behavior on compute module B's KVM.

The following example of the `sra_sysconfig` command helps you respond to the prompts given in subsequent steps in this procedure. Respond to the prompts with the information you wrote down earlier on the "Password and Installation Settings Printable Worksheet" on page 158. Text in bold font provides examples of the responses that you type.

Notes:

1. The example shows a static IP installation; a DHCP installation would prompt for fewer values.

2. There may be a pause of up to one minute during the EULA and Service Terms and Conditions prompts. Do not press any keys during this time.

```
2025/05/15 18:07:29 Welcome to the Stratus ztC Endurance System configuration Wizard
You must read and accept the Stratus End User License Agreement to continue.
Press Enter to continue to the Stratus End User License Agreement.

Do you accept the presented Stratus End User License Agreement? (yes/no): yes

You must read and accept the ztC Endurance Service Terms and Conditions to continue.
Press Enter to continue to the ztC Endurance Service Terms and Conditions.

Do you accept the presented terms and conditions? (yes/no): yes
Enter new BMC Web console admin password: zenadmin
Confirm new BMC Web console admin password: zenadmin
Enter new Management zenadmin password: zenadmin
Confirm new Management zenadmin password: zenadmin
2025/05/15 18:07:40 sra_sysconfig: Attempting to run Sysconfig script on standby...

[========================================] Sysconfig script completed successfully

Please enter Asset ID (eg <zenXXXXX>):  zen12345
Support notifications of system health information over the Active service network (ASN):
Enable support notifications [True/False]:                      False
ztC Endurance System Information:
Enter BMC A hostname:                                           sys20bmca
Enter BMC B hostname:                                           sys20bmcb
Select BMC Network protocol (dhcp/static):                      static
BMC A eth0 IP:                                                  10.109.30.229
BMC A eth1 IP:                                                  10.109.30.230
BMC B eth0 IP:                                                  10.109.30.231
BMC B eth1 IP:                                                  10.109.30.232
Enter Standby A Hostname:                                       sys20a
Enter Standby B Hostname:                                       sys20b
Select Standby Network protocol (dhcp/static):                  static
Standby B IP Address 1:                                         10.109.30.233
Standby B IP Address 2:                                         10.109.30.234
Standby A IP Address 1:                                         10.109.30.235
Standby A IP Address 2:                                         10.109.30.236
Standby Nameserver:                                             134.111.18.9
Enter RHEL Hostname:                                            sys20
Select RHEL Network Protocol (dhcp/static):                     static
RHEL IP Address:                                                10.109.30.237
RHEL Netmask:                                                   255.255.240.0
RHEL Gateway:                                                   10.109.16.1
RHEL Nameserver:                                                134.111.18.9
Enter RHEL Management Network Protocol (dhcp/static):           static
RHEL Management IP Address 2:                                   10.109.30.238
RHEL Management IP Address 1:                                   10.109.30.239
ztC Endurance Management Network Configuration  (BMC/Standby/Management):
Zen Management Gateway Address:                                 10.109.16.1
Zen Management Netmask:                                         255.255.240.0
Confirm Input below:
```

```
BMC:
bmca_hostname:                                          sys20bmca
bmcb_hostname:                                          sys20bmcb
bmc_netproto:                                           static
bmca_eth0:                                              10.109.30.229
bmca_eth1:                                              10.109.30.230
bmcb_eth0:                                              10.109.30.231
bmcb_eth1:                                              10.109.30.232

Standby:
standbya_hostname:                                      sys20a
standbyb_hostname:                                      sys20b
standby_netproto:                                       static
standbya_ipaddr_1:                                      10.109.30.235
standbya_ipaddr_2:                                      10.109.30.236
standbyb_ipaddr_1:                                      10.109.30.233
standbyb_ipaddr_2:                                      10.109.30.234
standby_nameserver:                                     134.111.18.9

Network (BMC, Standby and Host Mgmt):
zenmgmt_gateway:                                        10.109.16.1
zenmgmt_netmask:                                        255.255.240.0


RHEL:
rhel_hostname:                                          sys20
rhel_netproto:                                          static
rhel_ipaddr:                                            10.109.30.237
rhel_gateway:                                           10.106.16.1
rhel_nameserver:                                        134.111.18.9
rhel_netmask:                                           255.255.240.0

RHEL Management:
rhel_mgmt_netproto:                                     static
rhel_mgmt_ipaddr1:                                      10.109.30.239
rhel_mgmt_ipaddr2:                                      10.109.30.238
ASN ENABLED:                                            False
Asset ID:                                               zen12345
Do you want to continue with this configuration? [Yes/No]:        Yes
2025/05/15 18:10:57 Configuration saved successfully.
2025/05/15 18:10:57 Starting Mgmt Service..

Waiting for single node config to complete
[=======================================] Single-node config completed successfully
Running Zen Teaming script
[=======================================] Zen Teaming completed successfully
Script Output: Connection bond0 already exists


Performing Smart Exchange in 5 seconds..
2025/05/15 18:15:09 SetMaintModeCmd /api/inventory/compute/A/maintenance_mode?enable=true
requested maintenance mode on compute module A successfully
Waiting for Smart Exchange to complete

[=======================================] Smart Exchange completed successfully
2025/05/15 18:17:36 Peer online
Turning off maintenance mode for compute module: A
2025/05/15 18:17:36 SetMaintModeCmd /api/inventory/compute/A/maintenance_mode?enable=false
exited maintenance mode on compute module A successfully

[=======================================] Single-node config completed successfully
2025/05/15 18:21:38 Management Addresses for ztC Endurance console
```

```
2025/05/15 18:21:38 10.109.30.239
2025/05/15 18:21:38 10.109.30.238

Note: Host system requires a reboot to change the hostname.
Stratus ztC Endurance System configuration completed successfully. Press Enter to exit
```

c. At the `Enter the new BMC Web console admin password` prompt, type and confirm the new BMC Web console admin password. Press `Enter`.

d. At the `Enter the new Management zenadmin password` prompt, type and confirm the new `zenadmin` password. Press `Enter`.

   The `sra_sysconfig` script runs some scripts on compute module B, but you do not need to do anything on compute module B. If the script executes successfully, it displays the `Standby script executed successfully.` message on compute module A.

e. Type the asset ID for your system and press `Enter`.

f. Choose if you want to enable support notifications over the ASN. Type `True` or `False` and press `Enter`.

g. Type the two BMC hostnames. Type each entry and press `Enter`.

h. Type the BMC network protocol, `dhcp` or `static`, and press `Enter`.

i. If you selected `dhcp`, skip to the next step. Otherwise, if you selected `static`, type the four BMC IP addresses and press `Enter`.

j. Type the Standby hostnames for each compute module. Type each entry and press `Enter`.

k. Type the Standby network protocol, `dhcp` or `static`, and press `Enter`.

l. If you selected `dhcp`, skip to the next step. Otherwise, if you selected `static`, type the four Standby IP addresses and name server, if applicable. Type each entry and press `Enter`.

m. Type the RHEL hostname and press `Enter`.

n. Type the RHEL network protocol, `dhcp` or `static`, and press `Enter`.

o. If you selected `dhcp`, skip to the next step. Otherwise, if you selected `static`, type the RHEL IP address, gateway, netmask, and name server, if applicable. Type each entry and press `Enter`.

p. Type the RHEL management network protocol, `dhcp` or `static`, and press `Enter`.

q. If you selected `dhcp`, skip to the next step. Otherwise, if you selected `static`, type both RHEL management IP addresses. Type each entry and press `Enter`.

r. Type the Zen management network gateway address and press `Enter`.

s. Type the Zen management netmask address and press `Enter`.

t. After you have responded to all of the prompts, the system starts to configure both compute modules, and information about the configuration appears on the screen.

At the end of this process, the two IP addresses for connecting to the ztC Endurance console appear on the screen. **Make a note of these IP addresses for later use.**

> **Note**: If the system displays a message that the installation has failed to obtain an IP address, verify your DHCP server configuration after the installation has completed.

5. When you see the message `Note: Host system requires a reboot to change the hostname.`, reboot the system.

6. When you see the message `Stratus ztC Endurance System configuration completed successfully, Press Enter to exit`, press `Enter`.

Next Steps

Configure your system as described in "Post-Installation Tasks" on page 183.

Performing a Full System Software Installation

If Stratus did not install the system software at the factory, perform a full system software installation on your ztC Endurance system:

"I. Download a RHEL ISO Image" on page 167

"II. Power On the ztC Endurance System and Log On" on page 168

"III. Create the Installation Configuration File (Optional)" on page 169

"IV. Transfer the RHEL ISO Image to the ztC Endurance System" on page 173

"V. Configure BMC Network Settings" on page 175

"VI. Installing RHEL" on page 178

> **Note**: If you plan to customize your RHEL installation by adding configuring options to the kickstart file, see "Customizing the Installation Using the Kickstart File" on page 182 before proceeding with these instructions.

> **Notice**: Do not perform the following procedures until you have performed all of the procedures described in "Preparing to Install the System Software" on page 144.

I. Download a RHEL ISO Image

Overview

Before starting the system software installation, download the RHEL ISO image.

> **Notices**:
>
> 1. Download the RHEL ISO image only as described in this topic.
>
> 2. The USB storage device should meet the following requirements:
>
>   - It should be a known brand of USB storage device to ensure a successful software installation. For example, Stratus has successfully tested software installation with SanDisk$^{®}$ brand devices.
>
>   - It should be a new USB storage device to avoid introducing malware to your system.
>
>   - It should be formatted as NTFS and should have at least 16 GB of memory.

To download a RHEL ISO image from the Penguin Solutions Customer Service Portal

1. From a remote system, click the following link to access the **Penguin Solutions Customer Service Portal**:

   https://service.stratus.com

2. Click **Login**, enter your credentials, and click **Sign In**.

3. On the **Welcome to Customer Service** page, click **Software Downloads**.

4. On the **Software Downloads** page, select your product, operating system, and model. For example, if you want to download the ISO image for a ztC Endurance system running Automated Uptime Layer with Smart Exchange for RHEL, select the following values:

   a. Under **Product**, select **ztC Endurance**.

   b. Under **Operating System**, select **Linux**.

   c. Under **Model**, select the model you want.

ztC Endurance - Linux Downloads is displayed.

5. Click **ISO Image**.

6. Click **System Software Build RH-2.0.0.0 for RHEL 9.4 ISO** to download the ISO image.

7. After the download completes, use Linux commands to copy the RHEL ISO image to an NTFS-formatted USB storage device. This device is not intended to be bootable.

II. Power On the ztC Endurance System and Log On

To power on the system and log on

1. If the ztC Endurance system is already powered on, go to step 2.

   If the ztC Endurance system is **not** already powered on:

   a. Turn on the monitor(s).

   b. Press the power button on the control-panel ear of the system.

2. Use the keyboard connected to compute module B to log on to the Standby OS as the `zenadmin` user with the default password `zenadmin` (or other credentials, if provided).

III. Create the Installation Configuration File (Optional)

Optionally (but **recommended**), create an installation configuration file, `zen_config.yaml`, that contains the configuration settings that you gathered.

Preparing the Installation Configuration File

You can optionally create an installation configuration file that contains all of the system settings you wrote down on the "Password and Installation Settings Printable Worksheet" on page 158. (Note that the installation configuration file does not contain your passwords; you are prompted for them on the command line.) This installation configuration file, `zen_config.yaml`, is a template file located in `/opt/stratus/install` that you modify for your specific configuration. This file allows you to enter all of the configuration entries before starting the BMC configuration or software installation, thus avoiding typing mistakes during the installation. This file is also useful if you intend to install more than one system with similar parameters.

You use the same installation configuration file as input to the BMC configuration script (`bmc_config.sh`), software installation script (`install.sh`), or network configuration script (`net_config.sh`). Using this file means that you do not need to manually enter configuration settings while those scripts run.

If you exit out of the software installation script before completing the installation process, the system automatically saves a copy of your latest configuration entries in a configuration file with a time-stamped file name in the `/opt/stratus/install` directory.

> **Notes**:
>
> 1. If you use an installation configuration file and leave some settings blank, the installation script prompts you on the command line for those settings.
>
> 2. In the installation configuration file, if you inadvertently enter multiple values for the same setting, the installation script uses the last setting listed in the file.

Example of the Installation Configuration File

An example of the `zen_config.yaml` template file follows.

```
# SYSTEM INFORMATION
system_id:                              # Use '$SYSTEM-ID$' to include in hostname con-
struction
system_domain:                          # Use '$SYSTEM-DOMAIN$' to include in hostname
construction
system_nameserver:                      # Use '$NAMESERVER$ to use this for nameserver fields
below or '--NONE--'

# Information shared by all components on the 1G ztC Endurance management network
# BMCs, Standby
zenmgmt_gateway:                        # IPV4/V6 address or None (for dhcp)
zenmgmt_netmask:                        # IPV4 subnet address or None (for dhcp/ipv6/(if pre-
fix supplied in ip))
zenmgmt_ipv6_prefix:                    # IPV6 prefix or None (for dhcp/ipv4/(if prefix sup-
plied in ip))

# BMC
bmca_hostname: $SYSTEM-ID$bmca.$SYSTEM-DOMAIN$   # BMCA hostname (may be a fully qualified domain name)
bmcb_hostname: $SYSTEM-ID$bmcb.$SYSTEM-DOMAIN$   # BMCB hostname (may be a fully qualified domain name)
bmc_netproto:                           # BMC IP protocol ('dhcp' or 'static')
bmca_eth0:                              # BMC A eth 0 IPV4/V6 address(ipaddr/prefix) or None
(for dhcp)
bmca_eth1:                              # BMC A eth 1 IPV4/V6 address(ipaddr/prefix) or None
(for dhcp)
bmcb_eth0:                              # BMC B eth 0 IPV4/V6 address(ipaddr/prefix) or None
(for dhcp)
bmcb_eth1:                              # BMC B eth 1 IPV4/V6 address(ipaddr/prefix) or None
(for dhcp)

# Standby
standbya_hostname:  $SYSTEM-ID$a.$SYSTEM-DOMAIN$ # Standby A node hostname (may be a fully qualified
domain name)
standbyb_hostname:  $SYSTEM-ID$b.$SYSTEM-DOMAIN$ # Standby B node hostname (may be a fully qualified
domain name)
standby_netproto:                       # Standby IP protocol ('dhcp' or 'static')
standbya_ipaddr_1:                      # Standby IPV4/V6 address(ipaddr/prefix) or None (for
dhcp) for A eno1
standbya_ipaddr_2:                      # Standby IPV4/V6 address(ipaddr/prefix) or None (for
dhcp) for A eno2
standbyb_ipaddr_1:                      # Standby IPV4/V6 address(ipaddr/prefix) or None (for
dhcp) for B eno1
standbyb_ipaddr_2:                      # Standby IPV4/V6 address(ipaddr/prefix) or None (for
dhcp) for B eno2
standby_nameserver: $NAMESERVER$        # Standby IPV4/V6 address or '--NONE--' (for dhcp)

# RHEL
rhel_hostname: $SYSTEM-ID$.$SYSTEM-DOMAIN$       # RHEL hostname (may be a fully qualified domain
name)
rhel_bootproto:                         # IP protocol ('dhcp' or 'static')
rhel_ipaddr:                            # IPV4/V6 address(ipaddr/prefix) or None (for dhcp)
rhel_gateway:                           # IPV4/V6 address or None (for dhcp)
rhel_netmask:                           # IPV4 subnet address or None (for dhcp/ipv6/(if pre-
fix supplied in ip))
rhel_ipv6_prefix:                       # IPV6 prefix or None (for dhcp/ipv4/(if prefix
```

```
supplied in ip))
rhel_nameserver: $NAMESERVER$                         # IPV4/V6 address or '--NONE--'
rhel_iso:                                             # RHEL ISO path/URL

# RHEL Management (ztC Endurance UI console)
rhel_mgmt_netproto:  dhcp                             # RHEL Management network protocol ('dhcp' or
'static')
rhel_mgmt_ipaddr1:                                    # RHEL Management IPV4/V6 address(ipaddr/prefix) or
None (for dhcp)
rhel_mgmt_ipaddr2:                                    # RHEL Management IPV4/V6 address(ipaddr/prefix) or
None (for dhcp)

# ASN
asn_enabled:                                          # Indicates whether Support Notifications over ASN are
enabled ('yes' or 'no')
asn_asset_id:                                         # Asset Identifier

# EOF
```

Creating the Installation Configuration File

The template installation configuration file is located in the Standby OS, which runs Ubuntu. Therefore, you need to use a Linux-based text editor, such as `nano` or `vi`, to modify the file.

> **Notes**:
>
> 1. To use privileged commands in the `zenadmin` account, you need to add the `sudo` directive. If you open a file in `nano` without `sudo`, you cannot save the configuration file in its original, privileged directory.
>
> 2. When the Standby OS boots, you are in the `zenadmin` home directory, `/home/zenadmin`. To modify the `zen_config.yaml` file, you must change to the directory where the file is located: `/opt/stratus/install`.

Additional considerations:

- If you forgot to use the `sudo` directive, and the system does not allow you to save the configuration file, you can save your changes to a file in the `/home/zenadmin` directory. Later, you can specify the configuration file's path name on the command line, or you can simply specify its file name if `/home/zenadmin` is the current working directory.

- When typing entries in the configuration file, you must add a space between the colon (:) and the entry text (for example, `bmca_eth0: 10.10.40.32`).

- When you are typing and moving around the file, be careful not to insert any unsupported characters or text that would prevent the installation script from properly reading the file.

- Avoid modifying the configuration file in other operating systems or word processors that might add unsupported line feeds or special characters. Use only a plain text editor.

IV. Transfer the RHEL ISO Image to the ztC Endurance System

To transfer the RHEL ISO image to the ztC Endurance system with a USB storage device

1. Mount a USB storage device in the Standby OS on compute module B:

   a. Using the keyboard connected to compute module B, log on to the Standby OS as the `zenadmin` user with the default password `zenadmin` (or other credentials, if provided).

   b. Connect the USB storage device to one of the USB 3.0 ports on compute module B. To locate the USB 3.0 ports, see Figure 3-11.

   > **Notice**: Do not connect the USB storage device to the USB 2.0 port on the control-panel ear.

   c. Run the `transfer_iso.sh` command.

   > **Note**: To use privileged commands in the `zenadmin` account, you need to add the `sudo` directive.

   ```
   $ sudo /opt/stratus/bin/transfer_iso.sh


   [ <<< Locating the RHEL install ISO on current USB storage
   devices >>> ]

   Valid RHEL ISO found in device /mnt/zeniso/rhel-9.4-x86_64-
   dvd.iso
   -> Copying it to /home/zenadmin (this may take a while)

   #########################################################-
   ##############
   RHEL ISO copied successfully!
   ##
   Done. Exiting script!
   ```

d. To verify that the ISO transferred successfully, run the `ls` command after the script has exited. In the following example, the ISO appears in the command output.

```
$ ls /home/zenadmin
rhel-9.4-x86_64-dvd.iso
```

e. Physically remove the USB storage device.

V. Configure BMC Network Settings

Overview

Run the Baseboard Management Controller (BMC) configuration script, `bmc_config.sh`, as described in "To configure the BMC network settings" on page 176. Running the BMC configuration script configures the network settings for the BMCs in the ztC Endurance system. You can also configure the BMC network settings in the system BIOS configuration utility, but the BMC configuration script simplifies the changes and allows you to change the configuration on a running system.

Notices:

1. If you do not run the BMC configuration script, the system automatically uses DHCP settings for the BMC network addresses and retains the default BMC Web console password.

2. If you choose to run the BMC configuration script post-installation, you will also need to run the network configuration script (`net_config.sh`).

3. The software installation script (`install.sh`) that you run in "VI. Installing RHEL" on page 178 prompts for changes to the same BMC settings. However, you should modify the BMC settings with the BMC configuration script to prevent network disruptions in case the installation is being run over the BMC Web console. If you run the BMC configuration script now, type n to skip this step during the software installation.

Notes:

1. If you have the BMC Web console open for either compute module when you change network settings with the BMC configuration script, you will lose your connections. You need to reconnect to the new addresses you entered in the script or configuration file.

2. If you need to modify DDNS settings, use the **DNS Configuration** page of the BMC Web Console.

The BMC configuration script is located in `/opt/stratus/bin/bmc_config.sh`.

BMC configuration script usage:

```
Usage : bmc_config.sh {options}

Options:
-h            prints this menu
-u url        url with ISO image location
-f filename   path to ISO image
-y filename   path to configuration yaml file
```

Run the BMC configuration script with the installation configuration file (`zen_config.yaml`) to avoid manually entering the configuration settings when running the script.

> **Note**: If you run the `bmc_config.sh` script with the `zen_config.yaml` file, the BMC configuration settings are saved in a new `.yaml` file with a date and time stamp (for example, `config_26042023_180016.yaml`). You can then use the new `.yaml` file with the software installation script (`install.sh`).

To configure the BMC network settings

1. Start the BMC configuration script by running **one** of the following commands:

   > **Note**: To use privileged commands in the `zenadmin` account, you need to add the `sudo` directive.

   - If you created the installation configuration file: `sudo /opt/stratus/bin/bmc_config.sh -y /opt/stratus/install/zen_config.yaml`

     (In the preceding command, if you installed `zen_config.yaml` in a different directory, specify that directory instead.)

   - If you did not create the installation configuration file and want to respond to prompts for each setting: `sudo /opt/stratus/bin/bmc_config.sh {options}`

2. The BMC configuration script loads any existing settings from the installation configuration file and begins prompting for input.

> **Note**: If you already entered the settings into the installation configuration file, the BMC configuration script prompts only for the settings you have not entered.

3. The BMC configuration script prompts for the BMC configuration settings. Type new values for the settings you need to change and press `Enter` to continue. To accept an existing value, leave the field blank and press `Enter`. Repeat until the script displays **Confirm input**.

4. The BMC configuration script displays a summary of the settings you entered and prompts you to apply the settings or prompt for each setting again to make any corrections. Do one of the following:

    -   To accept the settings that you entered and apply them to the BMCs, type `y` and press `Enter`.

    -   To display each prompt again with passwords hidden, type `n` and press `Enter`. For each prompt, type any changes and press `Enter` to continue.

    -   To display each prompt again with passwords in clear text, type `unhide`. For each prompt, type any changes and press `Enter` to continue.

5. The BMC configuration script begins to implement the BMC settings that you specified, and then exits.

6. View the output to confirm that the BMC configuration script completed successfully.

VI. Installing RHEL

Overview

After running the BMC configuration script, run the software installation script, `install.sh`, as described in "To install RHEL and Automated Uptime Layer with Smart Exchange" on page 179. Running the software installation script automatically:

- Installs RHEL

- Installs Automated Uptime Layer with Smart Exchange components

- Mirrors the pair of system disks

- Mirrors any additional disk pairs in the system, if present

The software installation script is located in `/opt/stratus/bin/install.sh`.

Software installation script usage:

```
Usage : install.sh {options}
Options:
-h            prints this menu
-u url        url with AUL-SE installer image location
-f filename   path to AUL-SE installer ISO image
-y filename   path to configuration yaml file
```

Run the software installation script with the installation configuration file (`zen_config.yaml`) to avoid manually entering the configuration settings when running the script.

If you exit out of the software installation script before completing the installation process, the system automatically saves a copy of your latest configuration entries in a configuration file with a time-stamped file name in the `/opt/stratus/install` directory.

The software installation takes approximately 60 minutes.

Installing RHEL and AUL-Smart Exchange

To install RHEL and Automated Uptime Layer with Smart Exchange

1. Start the software installation script by running **one** of the following commands:

   - If you created the installation configuration file: `sudo /opt/stratus/bin/install.sh -y /opt/stratus/install/zen_ config.yaml`

     (In the preceding command, if you installed `zen_config.yaml` in a different directory, specify that directory instead.)

   - If you did not create the installation configuration file and want to respond to prompts for each setting: `sudo /opt/stratus/bin/install.sh {options}`

   The software installation script loads any existing settings from the installation configuration file and begins prompting for input.

2. Read the ztC Endurance EULA. To accept the terms and continue, type `y` and `Enter` to continue.

3. Read the ztC Endurance Service Terms and Conditions. To accept the terms and continue, type `y` and `Enter`.

4. The software installation script runs. If you already entered the settings into the installation configuration file, the installation script prompts only for the settings you have not entered. For any settings that you need to enter, see the settings you wrote down on the "Password and Installation Settings Printable Worksheet" on page 158.

   a. When prompted to configure the BMC network interfaces:

      - If you have already run `bmc_config.sh`, type `n`, and press `Enter`.

      - If you have not already run `bmc_config.sh`, type `y` to configure the interfaces now or `n` to configure them later, and press `Enter`.

> **Note**: If you are connected to the console of the compute module through the BMC Web console as described in "Connecting to the BMC Web Console" on page 272, you will lose your connection if you change the BMC network settings. If applicable, consider configuring the BMC network settings at a later time as described in "To configure the BMC network settings" on page 176 or "Updating Network IP Configuration Settings" on page 210.

b. Type the new BMC Web console `admin` password, the ztC Endurance Management `zenadmin` password, and the RHEL Administrator (`root`) password. Type each password and press `Enter`.

c. If you chose to configure the BMC network interfaces now, type the BMC hostnames. Type each entry and press `Enter`.

d. Type the BMC network protocol, `dhcp` or `static`, and press `Enter`.

e. If you selected `dhcp`, skip to the next step. Otherwise, if you selected `static`, type the four BMC IP addresses. Type each entry and press `Enter`.

f. Type the Standby hostnames for each compute module. Type each entry and press `Enter`.

g. Type the Standby network protocol, `dhcp` or `static`, and press `Enter`.

h. If you selected `dhcp`, skip to the next step. Otherwise, if you selected `static`, type the four Standby IP addresses and name server, if applicable. Type each entry and press `Enter`.

i. Type the RHEL hostname and press `Enter`.

j. Type the RHEL network protocol, `dhcp` or `static`, and press `Enter`.

k. If you selected `dhcp`, skip to the next step. Otherwise, if you selected `static`, type the RHEL IP address, gateway, netmask, and name server, if applicable. Type each entry and press `Enter`.

> **Notice**: In `net_config.sh`, IPv6 can be used for the primary RHEL interface but cannot be used for the RHEL Management Network.

l. Type the file path or URL to the RHEL ISO and press Enter. For example, if you copied the RHEL ISO to the Standby OS as described earlier in this procedure, the file path would be similar to the following:

    /home/zenadmin/rhel-9.4-x86_64-dvd.iso

m. Type the RHEL management network protocol, dhcp or static, and press Enter.

n. If you selected dhcp, skip to the next step. Otherwise, if you selected static, type both RHEL management IP addresses. Type each entry and press Enter.

o. Choose if you want to enable support notifications over the ASN. Type y or n and press Enter.

    Activating Support Notifications over the ASN allows the ztC Endurance system to send health and status notifications to Penguin Solutions Customer Service over a secure Internet connection.

p. Type the asset ID for your system and press Enter.

5. The software installation script displays a summary of the settings you entered and prompts you to begin the installation or prompt for each setting again to make any corrections. Do one of the following:

    ▪ To accept the settings that you entered and begin the installation, type y and press Enter. **The installation begins immediately.**

    ▪ To display each prompt again with passwords hidden, type n and press Enter. For each prompt, type any changes and press Enter to continue.

    ▪ To display each prompt again with passwords in clear text, type u or unhide. For each prompt, type any changes and press Enter to continue.

6. After you confirm your input, the software installation script begins to implement the settings that you specified and verify that all system components are ready and properly configured for installation.

7. Read the RHEL EULA. To accept the terms and continue, type y and Enter.

    After you accept the RHEL EULA, the software installation script continues the implementation and verification process.

8. When prompted to continue installing RHEL, type y and Enter to proceed.

9. The software installation script automatically installs RHEL and begins configuring the system software. The script reboots compute module A twice while it completes the configuration steps, begins mirroring the pair of system disks, and then mirrors any other pairs of data disks in the system.

10. Upon successful installation of the system, the software installation script displays the current system settings as well as the two IP addresses for the ztC Endurance console and the single IP address for the RHEL host. If you have not already done so, make note of the IP addresses now to ensure that you have this information.

> **Note**: If the system displays a message that the installation has failed to obtain an IP address, verify your DHCP server configuration after the installation has completed.

The software installation script also displays the `Installation completed successfully.` message.

Next Steps

Configure your system as described in .

Customizing the Installation Using the Kickstart File

The ztC Endurance installation process uses a kickstart file to enable automated installations. You can customize this file to specify options such as system language settings and network configurations.

> **Cautions**:
>
> 1. Sections of the kickstart file that handle ztC Endurance system software and required configuration are clearly marked. Do not modify them.
>
> 2. Modify kickstart files only if you are proficient in Linux system administration, as incorrect changes can lead to system misconfiguration or data loss.

Although the ztC Endurance installation process validates the kickstart file before making modifications, it may not detect all potential issues. If you have questions, contact Penguin Solutions Customer Service.

For more information, refer to the RHEL documentation on kickstart files.

Modifying the Kickstart File

Before installing the RHEL host, make the necessary changes to the kickstart file from the Ubuntu Standby OS.

The kickstart file is located in the default home directory of the `zenadmin` user on the Standby OS (`/home/zenadmin/customer_ks.cfg`). Make your additions in the marked section at the end of the file (shown here in bold):

```
############################################################################-
##########

#
# End of defined kickstart file
#

# For convenience, site changes may be placed here

## END ##
```

Before you start the RHEL installation, ensure that the customized `customer_ks.cfg` file is located in `/home/zenadmin/`. If the installation process does not find it in the expected location, the installation may halt or install using the default settings.

> **Note**: Ensure that you run `install.sh` from the standby compute module that includes the modified kickstart file.

Run the installation as described in "Performing a Full System Software Installation" on page 166.

Validation occurs during system installation, and installation halts if the customized kickstart file introduces errors.

Post-Installation Tasks

After installing the system software, you must complete some post-installation tasks, including:

- "Verifying the System Software Installation" on page 184
- "Configuring the BMCs for Remote System Management" on page 186
- "Configuring Network Interfaces " on page 186
- "Registering the System with the RHSM" on page 187

Related Topics

- "System Software Administration (Linux)" on page 196

Verifying the System Software Installation

> **Notice**: When you start or restart the RHEL host, RHEL requires several minutes to start. If you just restarted your system, `zen_verify` and other utilities can display proper output after 5 to 10 minutes.

To confirm that the system is configured correctly, run the following command in the Standby OS:

`$ sudo /opt/stratus/bin/zen_verify.sh`

> **Note**: To use privileged commands in the `zenadmin` account, you need to add the `sudo` directive.

If your system was properly installed, the output appears similar to the following, where all of the tests have passed.

```
$ sudo /opt/stratus/bin/zen_verify.sh
Enter RHEL hostname or IP address:      ocean.abc.com
Enter RHEL Username:     zenadmin
Enter RHEL password:
RHEL authentication success
Host RHEL OS Version check                          [PASS]
Host RHEL Stratus packages check                    [PASS]
Host RHEL raid check                                [PASS]
Host RHEL storage check                             [PASS]
Host RHEL network bond check                        [PASS]
Host RHEL network check                             [PASS]
Standby Ubuntu OS Version check                     [PASS]
Standby Ubuntu debian check                         [PASS]
BMC Firmware version check                          [PASS]
BIOS Firmware version check                         [PASS]
SES Firmware version check                          [PASS]
CPLD Firmware version check                         [PASS]
Smart Exchange compatibility                        [PASS]
```

The test descriptions are as follows:

- Host RHEL OS Version check

  Verifies that the RHEL host is running the supported version of RHEL.

- Host RHEL Stratus packages check

  Verifies that Stratus-specific packages have been installed on the RHEL host, and that they are the correct versions.

- Host RHEL raid check

  Verifies that the RHEL host has a proper RAID configuration.

- Host RHEL storage check

  Verifies that all disks are part of mirrored LUNs.

- Host RHEL network bond check

  Verifies that at least one of the embedded PCIe adapter ports is paired.

- Host RHEL network check

  Verifies that all network adapters are teamed based on symmetric slots in both I/O modules.

- Standby Ubuntu OS Version check

  Verifies that the ztC Endurance Standby OS is running the supported version of the Ubuntu OS.

- Standby Ubuntu debian check

  Verifies that Stratus-specific packages have been installed on the Standby OS, and that they are the correct versions.

- BMC Firmware version check

  Verifies that the system's baseboard management controller (BMC) firmware is up to date.

- BIOS Firmware version check

  Verifies that the system's Basic Input/Output System (BIOS) firmware is up to date.

- SES Firmware version check

  Verifies that the system's SCSI Enclosure Services (SES) firmware is up to date.

- CPLD Firmware version check

  Verifies that the system's Complex Programmable Logic Device (CPLD) firmware is up to date.

- Smart Exchange compatibility

Verifies that the compute modules are compatible for Smart Exchange. That is, the compute modules have the same BIOS version, the same `ce_mod` version, and the same amount of memory.

Configuring the BMCs for Remote System Management

Although you can manage some aspects of your ztC Endurance system from a monitor, keyboard, and mouse connected to the system (the local console), you typically manage all aspects of the Linux host from a remote management system.

Use the BMC Web console to perform lights-out management or open a connection to the local console of the ztC Endurance system from a remote management system. The BMC Web console allows you to monitor the state of your system, power the system on and off, manage the configuration of your Linux host, and view logs.

For more information about using the BMC Web console, see "Overview of the BMC Web Console" on page 271.

Configuring Network Interfaces

The installation process automatically configures the embedded 10GbE ports on both I/O modules (for example, `ethA0600` and `ethB0600`) into a channel-bonding interface, bond0.

> **Notice**: In the preceding bond0 interface, the MAC address of the bond can change. Read the manufacturer's documentation for your network switch and then configure the switch to allow MAC address changes.

Network devices are named according to the following scheme:

> eth*Module_letterSlot_numberPort_number*

Note that port 1 is represented by the value 1, and port 2 is represented by the value 0. For example, port 1 on a PCIe adapter in slot 5 on I/O module B would be named `ethB0501`, while port 2 on the same PCIe adapter would be named `ethB0500`.

The embedded 1GbE adapters always appear as `eno1` and `eno2`.

These channel-bonding interfaces are set to operate in active-backup mode (mode 1) with Dynamic Host Configuration Protocol (DHCP) enabled. In many cases, no additional configuration of these default bonds is necessary.

If your system contains non-embedded Ethernet PCIe adapters, you must also manually configure bonds for the additional ports. Place one port from each enclosure in each bonded pair.

The default ztC Endurance bonding mode is active-standby (also sometimes called Adapter Fault Tolerance, or AFT). While all Linux bonding modes are supported, only active-standby provides fault tolerance.

> **Note**: The TCP protocol recovers the packets because TCP resends the lost packets after its time-out period. Other network protocols, such as UDP and ICMP, do not resend lost packets.

To view all network devices and bonds, use the `ncmli dev` command as shown in the following example.

> **Note**: Use only NetworkManager (`nmcli`) to create new bonds with PCIe adapters.

```
[root@car90 input_decks]# nmcli dev
DEVICE     TYPE      STATE                  CONNECTION
bond0      bond      connected              bond0
bond3      bond      connected              bond3
eno1       ethernet  connected              eno1
eno2       ethernet  connected              eno2
ethA0200   ethernet  connected              ethA0200
ethA0600   ethernet  connected              bond-slave-ethA0600
ethB0200   ethernet  connected              ethB0200
ethB0600   ethernet  connected              bond-slave-ethB0600
lo         loopback  connected (externally) lo
ethA0100   ethernet  disconnected           --
ethA0101   ethernet  disconnected           --
ethA0201   ethernet  disconnected           --
ethA0601   ethernet  disconnected           --
ethB0100   ethernet  disconnected           --
ethB0101   ethernet  disconnected           --
ethB0201   ethernet  disconnected           --
ethB0601   ethernet  disconnected           --
```

Registering the System with the RHSM

Before you can obtain Red Hat software updates, you must register your system with Red Hat and have entitlement on the RHSM server. You must also limit downloads to specific versions.

- If you have not already registered, use the `subscription-manager-gui` command to open the Subscription Manager utility on the desktop, or use the `subscription-manager` command to enter subscription information directly on the command line. See your Red Hat documentation for further information.

- After registering your machine with the RHSM to receive updates, limit downloads to compatible minor versions only. For example, for systems running on Red Hat Enterprise Linux (RHEL) version 9.4, type the following command on the command line:

```
# subscription-manager release --set=9.4
```

Upgrading the System Software

This topic describes how to upgrade the Automated Uptime Layer with Smart Exchange system software on a Stratus ztC Endurance system. Upgrading the system can include upgrading the following:

- Red Hat Enterprise Linux (RHEL) and Stratus proprietary add-ins

- Compute module BMC firmware, BIOS firmware, and Standby OS software (Ubuntu)

> **Note**: Consider scheduling a maintenance window during which you can complete the upgrade process, which can take up to **3 hours**, depending on the number of components that require an upgrade. Because the system needs to upgrade software and firmware on both the A- and B-sides of the system, it must shut down the RHEL host halfway through the upgrade to perform a reboot exchange to the upgraded side and continue upgrading software and firmware on the other side.

Cautions:

1. Do not update the Ubuntu system software of the Standby OS from any source other than Stratus unless instructed to do so. To prevent incompatibilities and ensure the proper fault-tolerant operation of your system, use only the release and software packages that are installed with the Automated Uptime Layer with Smart Exchange software.

2. If you upgrade any software or firmware outside of the documented process for upgrades, the ztC Endurance system automatically reinstalls and reverts to the versions of software or firmware that are compatible with Automated Uptime Layer with Smart Exchange.

3. The system is not fault tolerant during an upgrade. The system will not be completely protected and able to perform a Smart Exchange until after the upgrade is complete and all modules are online.

The upgrade steps are:

"I. Review the Pre-Upgrade Checklist" on page 190

"II. Download the Upgrade File" on page 191

"III. Add the Upgrade File to the ztC Endurance console" on page 192

"IV. Upgrade the System Software" on page 193

"V. Complete Post-Upgrade Tasks" on page 196

I. Review the Pre-Upgrade Checklist

Review the following checklist before you upgrade the system software.

☐ For the latest information about upgrade paths and restrictions, see the <u>Release Notes: Automated Uptime Layer with Smart Exchange Release RH-2.0.0.0 (Linux)</u>.

☐ Prepare to accept the end-user license agreements (EULAs) that are presented during the upgrade process. To locate and read the text of the EULAs you must accept before upgrading Automated Uptime Layer with Smart Exchange, see <u>"End-User License Agreements and Warranty" on page 300</u>. You should not perform the upgrade if you cannot accept the EULAs or are not authorized to accept them. The upgrade process terminates if you decline any EULAs.

☐ Shut down any running applications. If applicable, confirm that your applications are configured to start automatically when the RHEL host starts. If you do not want your applications to start until after you have verified the success of the upgrade, you can disable automatic startup using Linux system and service management tools, but remember to re-enable it after the upgrade.

☐ With the ztC Endurance system halted, detach most peripheral devices from the system (exceptions are noted below). Devices to detach include:

- Unsupported PCIe adapters

- External storage systems

- USB devices, **except for** the keyboard, mouse, and USB storage device that you use to copy the upgrade kit

☐ Use the ztC Endurance console to verify that the system is in good health:

- From the **Hardware** page, check that all modules are running and not in maintenance mode.

- Resolve any alerts on the **Dashboard** page before continuing with the upgrade. The masthead must display a green check mark and **System is Healthy** message with no outstanding issues; otherwise, the upgrade will not proceed.

☐ Before scheduling a maintenance window, you can qualify your upgrade kit to ensure it is ready to apply. The **Qualify** button becomes active after you download an upgrade kit and add it to the ztC Endurance console **Upgrade Kits** page as described in the following sections.

☐   If you do not have an account on the **Penguin Solutions Customer Service Portal**, create one.

II. Download the Upgrade File

Downloading an upgrade kit requires access to the **Penguin Solutions Customer Service Portal**.

To download the Automated Uptime Layer with Smart Exchange upgrade kit

1. Click the following link to access the **Penguin Solutions Customer Service Portal**:

   https://service.stratus.com

2. Click **Login**, enter your credentials, and click **Sign In**.

3. On the **Welcome to Customer Service** page, click **Software Downloads**.

4. On the **Software Downloads** page, select the following values:

   a. Under **Product**, select **ztC Endurance**.

   b. Under **Operating System**, select **Linux**.

   c. Under **Model**, select the desired model.

   ztC Endurance – Linux Downloads is displayed.

5. At the bottom of the page, click **Upgrade Kit**.

   The page expands to show a description of the kit and the file to download, which is named similar to the following: **System Software Build 2.x.x.x for RHEL Upgrade Kit**.

6. Download the upgrade kit file.

III. Add the Upgrade File to the ztC Endurance console

After downloading the upgrade file to a local system, use the ztC Endurance console to add it to the **Upgrade Kits** page.

To add the kit to the Upgrade Kits page

1. Log in to the ztC Endurance console as the `admin` user.

2. In the ztC Endurance console, click **Upgrade Kits** in the left-hand navigation panel.

3. On the **Upgrade Kits** page, click the **Add** button beneath the masthead, which opens the **Upload Wizard**.

4. Browse to select and upload the file you downloaded from the customer service web site.

   The upgrade kit is added to the **Upgrade Kits** page.

5. After the upload is complete, close the wizard to return to the **Upgrade Kits** page.

   The **Upgrade Kits** page now shows the new upgrade kit, including its state and version number. The kit's state should be listed as **Ready**.

IV. Upgrade the System Software

For more information about the buttons on the **Upgrade Kits** page and the kit states listed in the **State** column, see "Upgrade Kits Page" on page 268.

To run the upgrade

1. On the **Upgrade Kits** page of the ztC Endurance console, select the upgrade kit to apply. To upgrade, the state of the kit must be shown as **Ready**.

2. Optionally, click the **Qualify** button and click **Yes** to confirm. The qualification process verifies that no version conflicts exist, and that the system is in a healthy state for the upgrade to start. If qualification fails, resolve the issue described in the error message.

3. Click **Upgrade**. The **Confirm Upgrade** alert appears.

4. Optionally, activate the **Automatically reboot system when required to complete upgrade** check box. This option allows the upgrade process to invoke automatic reboots as needed, without your confirmation.

5. Click **Yes** to start the upgrade.

   While an upgrade is in progress, the **Upgrade Kits** page lists the state of the **Host** (RHEL), **compute module A**, and **compute module B** as **Pending**, **Updating**, or **Up to date**.

   You can also monitor the progress of the upgrade through alerts displayed on the **Dashboard** page and by looking for **Autoburn Status** under the **Summary** tab on the **Hardware** page.

6. If you did not activate the **Automatically reboot system when required to complete upgrade** check box, you are prompted to reboot at some point during the upgrade. Confirm the reboot or follow the instructions for rebooting later. After rebooting, the following screens appear:

   a. The ztC Endurance console login screen. Do nothing.

   b. A window displaying **This site can't be reached**. Do nothing.

   When the ztC Endurance console appears again, log in. If the login fails, wait a few minutes until the system is fully booted and try again.

7. Continue to monitor the upgrade progress by opening the **Upgrade Kits** page.

Notice: The **Upgrade Kits** page notifies you if any upgrade step fails. If you encounter an error, contact Penguin Solutions Customer Service for assistance with clearing the error condition and issuing the `sysupdate - reset system` command. After the error is cleared, the upgrade resumes and displays the **Apply in progress** status again.

Caution: The **Upgrade Kits** page allows you to monitor the upgrade; however, there will still be long periods when it appears that there is no progress even while the upgrade continues in the background. **Do not interfere with the upgrade or restart the system; otherwise, you might leave the system in an unusable state.** If the upgrade takes longer than three hours with no signs of progress on the **Upgrade Kits** page, contact your authorized Penguin Solutions service representative for assistance before proceeding any further.

8. The system automatically completes the following tasks:

    ▪ Puts the standby compute module into maintenance mode

    ▪ Updates the ztC Endurance Red Hat Package Managers (RPMs) on the host

9. The ztC Endurance Management service (`zen-mgmt`) restarts after the RPMs are updated. This causes a log out of the ztC Endurance console. When `zen-mgmt` restarts, the host information reappears on the login screen.

    Before logging back in, refresh the browser to update any ztC Endurance console files the browser may have cached, and upload the updated ztC Endurance console files into the browser.

10. **If the kit does not contain a BIOS upgrade**, and you did not enable **Automatically reboot system when required to complete upgrade**, you are prompted to reboot the host. Confirm the reboot.

    After the host reboots, the upgrade continues. The system:

    ▪ Shuts down the Standby OS on the standby compute module

    ▪ Upgrades the BMC firmware on the standby compute module

    ▪ Upgrades the Standby OS software (Ubuntu) on the standby compute module

    ▪ Restarts the standby compute module

11. **If the kit contains a new BIOS**, the upgrade process triggers a reboot exchange, and waits for user acknowledgment before powering off the active compute module. If you enabled **Automatically reboot system when required to complete upgrade**, power-off occurs automatically.

    The RHEL host shuts down, which causes the upgraded standby to take over and boot the RHEL system on the compute module that is not powered down.

    > **Notes**:
    >
    > - You lose access to the ztC Endurance console for up to 20 minutes. After the RHEL host restarts, refresh your web browser and log in again.
    >
    > - After the reboot exchange, no further downtime is needed.

12. Wait for the exchange to complete. When the RHEL host is running on the newly upgraded compute module, software, and firmware, the system:

    - Puts the standby compute module into maintenance mode

    - Shuts down the Standby OS on the standby compute module

    - Upgrades the BMC and BIOS firmware on the standby compute module if necessary

    - Upgrades the Standby OS software (Ubuntu) on the standby compute module

    - Restarts the standby compute module

    - Takes the standby compute module out of maintenance mode

13. Continue monitoring the upgrade progress on the **Upgrade Kits** page of the ztC Endurance console.

    At this point, the **Upgrade Kits** page displays the state of the **Host** (RHEL) and the newly upgraded compute module as **Up to date**, and it displays the state of the standby compute module as **Updating packages**. You can also view the state of the compute modules on the **Dashboard** page.

After the system upgrades the standby compute module, the upgrade process is complete. On the **Upgrade Kits** page, the upgrade kit **State** changes to **Applied**.

V. Complete Post-Upgrade Tasks

1. In the ztC Endurance console, verify that the system is healthy with no outstanding alerts. Resolve and clear any alerts before continuing.

2. Optionally, delete the upgrade kit from the ztC Endurance console as described in the "Upgrade Kits Page" on page 268.

3. If the ztC Endurance console alerts you to register, you should register the system, download a license file, and apply the license file as described in "Registering the System" on page 239.

4. If you are using ASNConnect to allow service personnel to remotely diagnose your system, as described in "Enabling ASNConnect" on page 213, restart the ASN service with the following command:

```
sudo systemctl restart callhome
```

5. Verify that your system is functioning properly on the new software.

Chapter 6: System Software Administration (Linux)

Read the following topics to learn about tasks needed to administer a ztC Endurance system:

- "Managing Disks" on page 196
- "Updating Network IP Configuration Settings" on page 210
- "Configuring RHEL Components" on page 213
- "Enabling ASNConnect" on page 213
- "Installing a Custom Security Certificate" on page 214
- "Changing Passwords on a ztC Endurance System" on page 228

Managing Disks

Read the following topics to learn more about managing the disks on your ztC Endurance system:

- "Disk Storage Configuration" on page 197
- "Setting Up RAID Arrays on Provisioned Disks " on page 201
- "AUL-Smart Exchange Management of Failed Disks " on page 208
- "Secure Erase a Disk to Use as a Replacement Disk" on page 209

Disk Storage Configuration

The following sections describe disk naming and configuration for internal storage on a ztC Endurance system.

- "Internal Disk Naming " on page 197

- "Default Provisioned Disk Configuration for a Newly Installed System" on page 199

- "Managing Partitions" on page 200

Internal Disk Naming

A ztC Endurance system supports up to six disks in each storage module, for a total of 12 disks in each system. As on many Linux-based systems, the operating system assigns NVMe device names (for example, `nvme2` and `nvme4`) to each disk based on their order of discovery at system startup. However, the configuration may change dynamically when hardware failures occur, repairs are made, or when an administrator adds or removes components.

To simplify disk management, Automated Uptime Layer with Smart Exchange also assigns each disk a persistent user-device (`udev`) name based on its physical location in the system. Each disk name is associated with the slot, not the disk (for example, `diskA1`).

The disk names are developed according to the following scheme:

- The prefix `disk`

- A or B, indicating the storage module in which the disk resides

- 1, 2, 3, 4, 5, or 6, indicating the slot number

Figure 6-2 shows the persistent `udev` names for these disk slots in a fully-populated system.

Figure 6-2: udev Disk Names
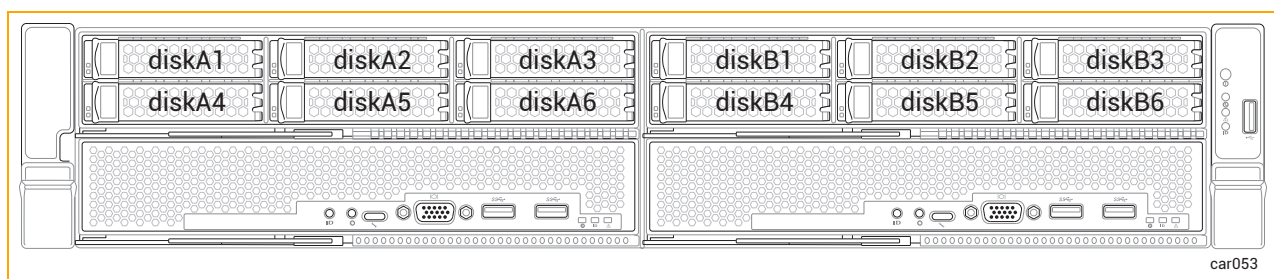
> **Note**: Although you can use the persistent udev names with standard Linux commands, only kernel device names (that is, NVMe disk names) appear in some files, such as `/var/log/messages`. Because the kernel device names are assigned on a first-come/first-served basis and are not persistent across disk swaps and reboots, determining which physical device a kernel device name refers to can be difficult. To translate kernel device names to their persistent udev names, use the udevadm command. For example, execute `udevadm info /dev/disk/by-dpid/diskA1` to display the names associated with the boot disk in slot 1 module A. Or, use the `ls` command to list all device names in both formats as described in the following example.

To display the persistent udev device names of all disks currently installed in the system, use the `ls` command to list the files in the `/dev/disk/by-dpid` directory. The following example shows sample output of the by-dpid directory:

```
ls -l /dev/disk/by-dpid/
total 0
lrwxrwxrwx. 1 root root 13 Apr  8 16:19 diskA1 -> ../../nvme0n1
lrwxrwxrwx. 1 root root 15 Apr  8 16:19 diskA1-part1 -> ../../nvme0n1p1
lrwxrwxrwx. 1 root root 15 Apr  8 16:19 diskA1-part2 -> ../../nvme0n1p2
lrwxrwxrwx. 1 root root 15 Apr  8 16:19 diskA1-part3 -> ../../nvme0n1p3
lrwxrwxrwx. 1 root root 15 Apr  8 16:19 diskA1-part4 -> ../../nvme0n1p4
lrwxrwxrwx. 1 root root 15 Apr  8 16:19 diskA1-part5 -> ../../nvme0n1p5
lrwxrwxrwx. 1 root root 13 Apr  8 16:19 diskA6 -> ../../nvme2n1
lrwxrwxrwx. 1 root root 15 Apr  8 16:19 diskA6-part1 -> ../../nvme2n1p1
lrwxrwxrwx. 1 root root 13 Apr  8 16:19 diskB1 -> ../../nvme4n1
lrwxrwxrwx. 1 root root 15 Apr  8 16:19 diskB1-part1 -> ../../nvme4n1p1
lrwxrwxrwx. 1 root root 15 Apr  8 16:19 diskB1-part2 -> ../../nvme4n1p2
lrwxrwxrwx. 1 root root 15 Apr  8 16:19 diskB1-part3 -> ../../nvme4n1p3
lrwxrwxrwx. 1 root root 15 Apr  8 16:19 diskB1-part4 -> ../../nvme4n1p4
lrwxrwxrwx. 1 root root 15 Apr  8 16:19 diskB1-part5 -> ../../nvme4n1p5
lrwxrwxrwx. 1 root root 13 Apr  8 16:19 diskB6 -> ../../nvme6n1
lrwxrwxrwx. 1 root root 15 Apr  8 16:19 diskB6-part1 -> ../../nvme6n1p1
```

The output also indicates that the pair of boot disks, `diskA1` and `diskB1`, have five partitions (for example, `disknnnnnn-partn`), whereas a second pair of disks, `diskA6` and `diskB6`, have one partition.

If you execute the `ls` command with the `-l` option (to display the long-listing format, as shown), you can see that the persistent `disknnnnnn` names are symbolic links that point back to the original NVMe device names assigned by the operating system. This is helpful information if you find references to the

NVMe device names in files or command output and you need to use the `disknnnnnn` name to determine the location of the disk in the system.

You can specify `udev` names in either format with standard Linux commands.

For fault tolerance, the disks are paired based on like slots of disks in the storage modules: `diskA1` is paired with `diskB1`, `diskA2` with `diskB2`, and so on. RAID 1 arrays are created by placing one mirror on each disk of the pair. For example, RAID array `/dev/boot` occupies partitions on `diskA1` and `diskA2`.

Default Provisioned Disk Configuration for a Newly Installed System

By default, Linux is installed on disks in slot 1 of storage module A and storage module B.

During the installation, Linux dynamically assigns `nvme` device names to these disks (for example, `nvme0n1` and `nvme1n1`). After Automated Uptime Layer with Smart Exchange is fully installed, it automatically assigns persistent device names beginning with labels that identify storage module A or B and slot number 1 through 6. For example, provisioned NVMe disks in slot 1 of storage module A and B would be named `diskA1` and `diskB1`, respectively.

All assigned persistent names are listed under the device directory `/dev/disk/by-dpid`.

Table 6-4 describes the default partitions and RAID arrays on the system disks after Linux is installed. The table assumes that the system disks reside in slot 1 of storage module A and storage module B.

Table 6-4: Default System Disk Partitions

Mount Points	Size	RAID Array	Mirrored Partitions
`(swap)`	4 GB	`/dev/md/swap`	diskA1-part3 diskB1-part3
`/(root)`	69.9 GB	`/dev/md/root`	diskA1-part2 diskB1-part2
`/boot`	1022 MB	`/dev/md/boot`	diskA1-part4 diskB1-

Mount Points	Size	RAID Array	Mirrored Partitions
			part4
/boot/efi	599.9 MB	/dev/md/efi	diskA1-part5 diskB1-part5
/home	1.4 TB	/dev/md/home	diskA1-part1 diskB1-part1

> **Note**: Device names for mirrored partitions on internal disks may vary by system.

Managing Partitions

Use the `parted` utility to display and change a disk's partition table and geometry. For details, see `parted(8)` and [Getting started with partitions](#) in the Red Hat product documentation.

The following example demonstrates how to use the `parted` utility to partition a new provisioned disk, `diskB6`. If applicable, partition `diskA6` in the same manner before creating a RAID array as described in ["Creating RAID Arrays" on page 202](#). When adding and partitioning internal disks to be part of a new RAID pair, ensure that both disks have the same geometry, partition table, and type.

> **Notice**: Creating a new partition table destroys any existing data on the target disk.

```
# parted /dev/disk/by-dpid/diskB6
GNU Parted 3.5
Using /dev/nvme7n1
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted)  mkpart xfs 0% 100%
(parted) p
Model: MTFDKCC6T4TGQ-1BK1JABYY (nvme)
Disk /dev/nvme7n1: 6401GB
Sector size (logical/physical): 512B/4096B
```

```
Partition Table: gpt
Disk Flags:

Number  Start    End      Size     File system  Name  Flags
1       1049kB   6401GB   6401GB                xfs


(parted) q
Information: You may need to update /etc/fstab.



# ls -l /dev/disk/by-dpid/diskB6*
lrwxrwxrwx. 1 root root 13 Jun  3 15:22 /dev/disk/by-dpid/diskB6 -> ../../n-
vme7n1
lrwxrwxrwx. 1 root root 15 Jun  3 15:22 /dev/disk/by-dpid/diskB6-part1 ->
../../nvme7n1p1
# grep nvme7n1 /proc/partitions
259        26 6251223384 nvme7n1
259        28 6251222016 nvme7n1p1
```

Setting Up RAID Arrays on Provisioned Disks

This section discusses the following topics related to creating RAID arrays for the internal disk drives:

- "RAID Array Overview" on page 201

- "Creating RAID Arrays" on page 202

- "Creating and Mounting a File System" on page 205

- "Checking the Current State of RAID Arrays " on page 206

RAID Array Overview

RAID, or Redundant Array of Independent Disks, is the basis for fault-tolerant file system availability. As disks go in and out of service, the only way to keep the file system available is to mirror it on multiple disks, with a provisioned disk in each storage module.

All file systems are created on RAID devices. Automated Uptime Layer with Smart Exchange supports RAID 1 (mirrored) and RAID 10 (striped) on RAID 1 provisioned disks. You configure the RAID 10 array with RAID 1 devices, since the underlying devices must be fault tolerant.

The default format for device names for a RAID array consists of the letters md and a number from 0 through 127 (for example, /dev/md23). The number must be unique among the running RAID arrays.

> 🛈 **Notice**: Creating a RAID array destroys any existing data on the disk partitions.

> 🛈 **Note**: To ensure fault-tolerant operation of the system, use only the `mdadm` command to create RAID arrays, and only use device names in the default format.

Creating RAID Arrays

This section describes how to create RAID 1 arrays, as well as how to create RAID 10 (RAID 1 + 0) arrays by combining RAID 1 arrays.

> 🛈 **Note**: Although you can use the persistent `udev` names with standard Linux commands, some files display only the kernel device names, so the command output may appear different from what you entered. See "Internal Disk Naming " on page 197 for information on determining which kernel device name pairs with which `udev` name.

Creating a RAID 1 Array

The following procedure describes how to create a RAID 1 array.

To create a RAID 1 array

1. Select a pair of same-sized disks and insert them into two corresponding slots in different storage modules as described in "Removing and Inserting a Disk Drive" on page 113. This example uses the disk in the right-most slot of the bottom row of storage module A (`diskA6`) and storage module B (`diskB6`) as seen in Figure 1-8: Disk-Drive Slot Numbering.

2. When the disks have spun up, partition them for the desired RAID array, as described in "Managing Partitions" on page 200. Create partitions of equal size on each disk to be paired.

3. If the target disk partitions have existing RAID information, secure erase each target partition as described in "Secure Erase a Disk to Use as a Replacement Disk" on page 209 before proceeding.

4. Determine a device name for the new array (for example, `/dev/md20`). Verify that the device name is not already in use by issuing the following command:

```
# cat /proc/mdstat
Personalities : [raid1]
md98 : active raid1 nvme4n1p1[2] nvme0n1p1[3]
```

```
20953088 blocks super 1.2 [2/2] [UU]


md123 : active raid1 nvme3n1p5[1] nvme1n1p5[0]
614336 blocks super 1.0 [2/2] [UU]
bitmap: 0/1 pages [0KB], 65536KB chunk


md124 : active raid1 nvme3n1p1[0] nvme1n1p1[1]
1483422720 blocks super 1.2 [2/2] [UU]
bitmap: 0/12 pages [0KB], 65536KB chunk


md125 : active raid1 nvme1n1p3[0] nvme3n1p3[1]
4189184 blocks super 1.2 [2/2] [UU]


md126 : active raid1 nvme1n1p2[1] nvme3n1p2[0]
73333760 blocks super 1.2 [2/2] [UU]
bitmap: 1/1 pages [4KB], 65536KB chunk


md127 : active raid1 nvme1n1p4[1] nvme3n1p4[0]
1046528 blocks super 1.2 [2/2] [UU]
bitmap: 0/1 pages [0KB], 65536KB chunk
```

5. To create and start the RAID array, issue a command similar to the following, where the backslash (\) indicates line continuation:

```
# sudo mdadm -C /dev/md20 -b internal --level=1 --raid-devices=2 \
/dev/disk/by-dpid/diskA6-part1 \
/dev/disk/by-dpid/diskB6-part1
```

This command creates the RAID array and starts it. The -b internal option implements Red Hat Linux RAID 1 fast resync, which can greatly reduce the amount of time the arrays need to resynchronize (though the resynchronization process will likely still take multiple hours). To prevent system instability, use this option when you create the array, instead of using the --grow option at a later time.

To see the status of the new RAID array, type a command similar to the following:

```
# mdadm --detail /dev/md20
```

Creating a RAID 10 Array

When you want the file system to be larger than a RAID 1 volume, you can create a RAID 10 array, which consists of at least two RAID 1 arrays joined together with a RAID 0 array.

The following example assumes that two RAID 1 arrays have been created:

- md20, consisting of the partitions diskA3-part1 and diskB3-part1

- md21, consisting of the partitions diskA4-part1 and diskB4-part1

To create a RAID 10 array

1. Select a pair of RAID 1 arrays.

2. Use the following command to create and start a RAID 10 array:

   # **sudo mdadm --create /dev/md30 --level=0 --raid-devices=2 /dev/md20 /dev/md21**

3. To display the active RAID array, type the following command:

   # **sudo mdadm --detail /dev/md30**

   The output appears similar to the following example:

   ```
   # mdadm --create /dev/md30 --level=0 --raid-devices=2 /dev/md20 /dev/md21
   mdadm: Defaulting to version 1.2 metadata
   mdadm: array /dev/md30 started.
   # mdadm --detail /dev/md30
   /dev/md30:
   Version : 1.2
   Creation Time : Thu Jun 12 16:58:13 2025
   Raid Level : raid0
   Array Size : 4687899648 (4.37 TiB 4.80 TB)
   Raid Devices : 2
   Total Devices : 2
   Persistence : Superblock is persistent
   Update Time : Thu Jun 12 16:58:13 2025
   State : clean
   Active Devices : 2
   Working Devices : 2
   Failed Devices : 0
   Spare Devices : 0
   Layout : original
   Chunk Size : 512K
   Consistency Policy : none
   Name : sys04.ocean.abc.com:30  (local to host sys04.ocean.abc.com)
   UUID : e366b2a5:e06c1ed1:8b8dde70:46bf46a6
   ```

```
Events : 0
Number    Major    Minor    RaidDevice State
0          9        20        0       active sync    /dev/md20
1          9        21        1       active sync    /dev/md21
```

Creating and Mounting a File System

The RAID arrays you create by following the instructions in "Creating RAID Arrays" on page 202 are raw disk block devices. To use the file system, you make it a journaling file system and mount it.

1. Make the file system a journaling file system (for example, XFS) by typing a command similar to the following:

    # **mkfs.xfs /dev/md30**

2. Mount the file system on a convenient mount point (for example, /mnt/big_data) by typing commands similar to the following:

    # **mkdir /mnt/big_data**

    # **mount /dev/md30 /mnt/big_data**

3. Verify that the file system is mounted by doing one or both of the following:

    ▪ Type an ls command, similar to the following, which shows the lost+found directory in the file system:

      # **ls -l /mnt/big_data**

    ▪ Type a df command, similar to the following, which shows the size of the file system:

      # **df /mnt/big_data**

4. Retrieve the UUID from the filesystem on the md RAID device by typing commands similar to the following:

    # **blkid | grep md30**

    /dev/md30: UUID="8316b178-f756-4532-be1b-098d2648ffc0" TYPE="xfs"

5. Add the mount to the /etc/fstab file using the UUID retrieved in the previous step, so the file system is mounted each time the system is rebooted. For example, add a line similar to the following:

    UUID="8316b178-f756-4532-be1b-098d2648ffc0"/mnt/big_data xfs defaults 0 0

6. Reboot the system to verify that the file system is mounted.

> **Note**: For data fault-tolerance, do not use a single disk that is not part of a RAID 1 array. An alert is generated if you attempt to mount a partition on a single disk that is not part of a RAID 1 array.

Checking the Current State of RAID Arrays

The current states of RAID volumes and physical RAID members can be monitored by logging into the ztC Endurance system and navigating to the "Hardware Page" on page 229 or the "Volumes Page" on page 238.

RAID volumes can also be monitored from a command line. Two sources provide information about the current state of RAID: the `/proc/mdstat` file and the `mdadm --detail` command.

The `/proc/mdstat` file provides an overview of active RAID devices, disk mirror status, and disk resynchronization progress, but it uses kernel device names that do not correspond to Automated Uptime Layer with Smart Exchange user-space device names. These names can cause confusion, and you cannot use them to identify particular disks.

The `/sbin/mdadm --detail` command uses user-space device names, enabling you to identify the disks described in the command output with particular physical disks. However, the `mdadm --detail --scan` command, which provides an overview of running mirrors, does not display disk resynchronization progress. To see resynchronization progress, you must issue `mdadm --detail` on a specific mirror. Therefore, use these sources as follows:

- Use the `/proc/mdstat` file when you need to check only the resynchronization progress on a particular mirror, since `/proc/mdstat` shows the resynchronization status for all mirrors.

- Use the `mdadm --detail /dev/md`*xxx* command to see the disk resynchronization status of a specific device (represented by *xxx*). Note that this command displays the disk's kernel device name, not the persistent `udev` name. To list all devices with their `udev` names and corresponding kernel device names, use `ls /dev/disk/by-dpid` as described in "Internal Disk Naming " on page 197.

The following example shows the contents of a sample `/proc/mdstat` file. The disk device names (nvme*) do not correspond to the names that Automated Uptime Layer with Smart Exchange installation assigns to internal disks based on their storage module and slot location.

The `/proc/mdstat` File

```
# cat /proc/mdstat
Personalities : [raid1]
md98 : active raid1 nvme4n1p1[2] nvme0n1p1[3]
```

```
20953088 blocks super 1.2 [2/2] [UU]


md123 : active raid1 nvme3n1p5[1] nvme1n1p5[0]

614336 blocks super 1.0 [2/2] [UU]

bitmap: 0/1 pages [0KB], 65536KB chunk


md124 : active raid1 nvme3n1p1[0] nvme1n1p1[1]

1483422720 blocks super 1.2 [2/2] [UU]

bitmap: 0/12 pages [0KB], 65536KB chunk


md125 : active raid1 nvme1n1p3[0] nvme3n1p3[1]

4189184 blocks super 1.2 [2/2] [UU]


md126 : active raid1 nvme1n1p2[1] nvme3n1p2[0]

73333760 blocks super 1.2 [2/2] [UU]

bitmap: 1/1 pages [4KB], 65536KB chunk


md127 : active raid1 nvme1n1p4[1] nvme3n1p4[0]

1046528 blocks super 1.2 [2/2] [UU]

bitmap: 0/1 pages [0KB], 65536KB chunk
```

The following example shows the use of the `mdadm` command to display the resynchronization progress of a particular device, `md51`:

The `mdadm --detail /dev/md`*`xxx`* Command

```
# mdadm --detail /dev/md51
/dev/md51:
Version : 1.2
Creation Time : Tue Apr  1 13:16:23 2025
Raid Level : raid1
Array Size : 1531419648 (1460.48 GiB 1568.17 GB)
Used Dev Size : 1531419648 (1460.48 GiB 1568.17 GB)
Raid Devices : 2
Total Devices : 2
Persistence : Superblock is persistent
Intent Bitmap : Internal
Update Time : Tue Jun  3 14:12:28 2025
State : clean, degraded, recovering
```

```
Active Devices : 1
Working Devices : 2
Failed Devices : 0
Spare Devices : 1
Consistency Policy : bitmap
Rebuild Status : 51% complete
Name : sys04.ocean.abc.com:51  (local to host sys04.ocean.abc.com)
UUID : 7a8e34de:34fcada0:84bfe34c:a733c1e4
Events : 150273
Number   Major   Minor   RaidDevice State
2     259       10        0        spare rebuilding   /dev/nvme2n1p1
1     259       21        1        active sync    /dev/nvme6n1p1
```

AUL-Smart Exchange Management of Failed Disks

When a disk that is part of a RAID set experiences a fault, the system determines whether the disk has experienced an unacceptable number of errors by calculating its mean time between failures (MTBF).

If the MTBF is currently greater than the assigned MTBF threshold, the system temporarily removes the disk from its RAID 1 mirror and then adds the disk back. If the MTBF drops below the threshold, the system removes the affected disk from its RAID mirror and removes the disk from service. If the system is configured to send alerts to the ASN, the system sends an alert.

If a disk has had an MTBF failure, MGMT faults the drive and removes it from service.

> **Notes**:
>
> 1. The last active disk of a RAID set is never removed from service, even if errors are reported against it.
>
> 2. The alert signals your authorized Penguin Solutions service representative to send a replacement disk.

Every time an active disk is pulled, all outstanding I/O is returned as errors, and all RAID members or mirrors that have active I/O on the missing disk are marked as faulty.

> **Notice**: Never remove the only remaining or healthy disk of a disk mirror. Doing so could cause a data loss or halt system operation.

When the system removes a disk from service, the disk of the remaining active mirror and the storage module that contains the remaining active disk enter a simplex state and are no longer safe to pull. A

blinking amber LED on the control-panel ear indicates that the disk drive is in a simplex state and is unsafe to remove. See for more information.

> **Notice**: Never remove a disk that is in a `SIMPLEX` state. Doing so could cause a data loss or halt system operation.

Secure Erase a Disk to Use as a Replacement Disk

If you are replacing a failed disk with a disk that is not blank, use the following procedures to manually verify that a spare disk is not in use, and, if applicable, perform a secure erase on the disk before using it as a replacement disk.

-

-

To verify that the spare disk is not in use

1. Check to see if the disk you inserted is already in use as a member of a RAID array by typing the following command:

   `# mdadm --detail --scan --verbose`

2. Check to see if the disk is in use to provide swap space by typing the following command:

   `# swapon -s`

3. Check to see if the disk is in use because it contains a mounted file system by typing the following command:

   `# cat /etc/mtab`

If the kernel device name for the disk is listed in the output of any of the preceding commands, the disk is in use, and you must choose another disk. Otherwise, secure erase the disk as described in .

The preceding commands display the kernel device name. See for information about using the `ls` command to map the kernel device name to the persistent `udev` name.

To perform a secure erase on a spare disk drive

After verifying that the spare disk is not in use, secure erase the drive to be added using the following command, in which *device_name* is the NVMe device you want to secure erase:

    sudo nvme format -s1 *device_name*

The *device_name* can be in the form `/dev/nvme1n1` or the udev equivalent form `/dev/disk/by-dpid/diskB6`.

If the disk you are attempting to secure erase is in use, a warning similar to the following appears:

```
# nvme format -s1 /dev/disk/by-dpid/diskB6
/dev/disk/by-dpid/diskB6: Device or resource busy
Failed to open diskB6.
Namespace is currently busy.
```

If the disk is not in use and the secure erase is successful, a message similar to the following appears:

```
# nvme format -s1 /dev/disk/by-dpid/diskB6
You are about to format diskB6, namespace 0x1.
WARNING: Format may irrevocably delete this device's data.
You have 10 seconds to press Ctrl-C to cancel this operation.
Sending format operation ...
Success formatting namespace:1
```

Updating Network IP Configuration Settings

Use the network configuration script (`net_config.sh`) if you need to update the network settings for the ztC Endurance system after the initial software installation. For example, you might need to change the network settings if you prepared a system with settings for an end user, but then you need to update the network settings to deploy the system on a different network or subnet, possibly at a new location.

The network configuration script updates network settings for the:

- Baseboard management controllers (BMCs)

- Standby OS of each compute module

- RHEL host

> **Notice**: When you run the network configuration script, it restarts the active compute module twice to complete all of the changes. Because you temporarily lose access to your applications, ensure that you run the script only on a non-production system or during a planned maintenance period.

> **Notes**:
>
> 1. If you are preparing for the initial software installation, use the BMC configuration script (`bmc_config.sh`) or the software installation script (`install.sh`) to configure the network settings instead of using the network configuration script. See the instructions in "III. Create the Installation Configuration File (Optional)" on page 169.
>
> 2. If you are connected to the console of the compute module through the BMC Web console, you will lose your connection if you change the BMC network settings. You need to reconnect to the new addresses you entered in the script or configuration file.

If you want to create a file that already contains the new network settings needed for the network configuration script, see "III. Create the Installation Configuration File (Optional)" on page 169. You can specify this file for input when running the network configuration script, BMC configuration script, or software installation script.

Network configuration script usage:

```
Usage : net_config.sh {options}


Options:
-h              prints this menu
-u url          url with ISO image location
-f filename     path to ISO image
-y filename     path to configuration yaml file
```

To update the system network settings

1. Optionally, create a file that already contains the configuration settings as described in "III. Create the Installation Configuration File (Optional)" on page 169; otherwise, be prepared to manually enter the information you wrote down on the "Password and Installation Settings Printable Worksheet" on page 158.

2. Log on to the Standby OS of either compute module as the `zenadmin` user.

3. Start the network configuration script by running **one** of the following commands:

   > **Note**: To use privileged commands in the `zenadmin` account, you need to add the `sudo` directive.

- If you created the configuration file: `sudo /opt/stratus/bin/net_config.sh -y /opt/stratus/install/zen_config.yaml`

  (In the preceding command, if you installed `zen_config.yaml` in a different directory, specify that directory instead.)

- If you did not create the configuration file and want to respond to prompts for each setting: `sudo /opt/stratus/bin/net_config.sh {options}`

4. The configuration script loads any existing settings from the installation configuration file and begins prompting for input.

> 🛈 **Note**: If you already entered the settings into the installation configuration file, the network configuration script prompts only for the settings you have not entered.

5. Type the current `zenadmin` password and press Enter to continue.

6. Type the RHEL Administrator password and press Enter to continue.

7. The script prompts for the IP configuration settings for the BMCs, the Standby OS, and the RHEL host. Type new values for the settings you need to change and press Enter to continue. To accept an existing value, leave the field blank and press Enter. Repeat until the script displays **Confirm input below**.

8. The script displays a summary of the settings you entered and prompts you to apply the settings or prompt for each setting again to make any corrections. Do one of the following:

   - To accept the settings that you entered and apply them to the system, type y and press Enter.

   - To display each prompt again with passwords hidden, type n and press Enter. For each prompt, type any changes and press Enter to continue.

   - To display each prompt again with passwords in clear text, type `unhide`. For each prompt, type any changes and press Enter to continue.

9. The configuration script begins to implement the network settings that you specified. View the output to confirm that the changes are successful.

10. The current login session to the Standby OS automatically disconnects so that the system can change the network settings for the compute module.

11. After the system restarts, verify that the BMCs, the Standby OS, and the RHEL host are functioning properly.

Related Topics

-

Configuring RHEL Components

After you finish performing the , you can configure the RHEL host.

If you need to configure additional network connections, you can configure PCIe adapter bonds as discussed in .

Enabling ASNConnect

ASNConnect by BeyondTrust is a third-party remote access solution that allows Penguin Solutions Customer Service to remotely connect to various components of a ztC Endurance system for assistance with diagnostics, troubleshooting, and service/support activities.

ASNConnect allows Penguin Solutions Customer Service to connect remotely to the following:

- On BMC A: the SSH interface and the BMC Web console

- On BMC B: the SSH interface and the BMC Web console

- On the RHEL host: the RHEL desktop, the file transfer interface, and the ztC Endurance console

- On the standby compute module: the SSH interface

To use ASNConnect, you must install software on your ztC Endurance system.

Checking Your Connection to ASNConnect

Before you install ASNConnect on your ztC Endurance system, use the `curl` command to make sure that your ztC Endurance system can connect to **asnconnect.stratus.com** over TCP port 443. The following example shows a successful connection.

```
curl -v asnconnect.stratus.com:443
* About to connect() to asnconnect.stratus.com port 443 (#0)
*   Trying 3.145.54.234...
* Connected to asnconnect.stratus.com (3.145.54.234) port 443 (#0)
> GET / HTTP/1.1
> User-Agent: curl/7.29.0
> Host: asnconnect.stratus.com:443
> Accept: */*
>
* Empty reply from server
* Connection #0 to host asnconnect.stratus.com left intact
```

Installing ASNConnect on the ztC Endurance System

> **Note**: You must complete the license registration for your system before attempting to install ASNConnect. See "Registering the System" on page 239 for more information.

1. Log in to the ztC Endurance console.

2. In the left-hand navigation panel, click **Support Configuration** under **Remote Support**.

3. Click the check box next to **Enable Remote Support Access**. This automatically installs ASNConnect and enables it on the Standby OS.

4. Click **Save** (or click **Cancel** to restore the previously saved values).

> **Note**: After the installation is complete, perform a Smart Exchange to auto-install and enable ASNConnect on the new Standby OS. See "Using Maintenance Mode" on page 235 for more information.

Your ztC Endurance system is now ready to use ASNConnect.

Related Topics

- "Post-Installation Tasks" on page 183

- "ztC Endurance Console" on page 218

- "Configuring Remote Support Settings" on page 264

Installing a Custom Security Certificate

The ztC Endurance management server automatically initializes itself with a self-signed SSL certificate, allowing communications to the management server to be encrypted without the identity of the certificate holder being verified by a third party. To view details of this certificate, click the warning icon on the left side of the browser's address bar.

If you prefer to install a custom certificate on the server, run the `certificate_installer` script. You can also use this script to revert to the default certificate or to display the current certificate.

The installation script is located in `/opt/stratus/sbin/` and must be run as `root`. The installer requires a key file containing the private key in addition to the certificate file. Both files must have the `.pem` format.

Installation script usage:

```
$ sudo /opt/stratus/sbin/certificate_installer.sh
To install/recover/display the certificate:

Usage: sudo /opt/stratus/sbin/certificate_installer.sh [Option]
-d, displays current certificate
-r, recovers to default certificate
-c [PATH to certificate.pem] -k [PATH to key.pem], installs the customer cer-
tificate
-h, displays help
$
```

> **Note**: The `certificate_installer` script automatically restarts the management server after any operation that changes the certificate, so you do not need to perform a manual restart after installation.

To install a custom certificate

1. Copy your new, valid certificate and key `.pem` files to the `/tmp` directory of the ztC Endurance RHEL host.

2. Issue the following command:

   ```
   $ sudo /opt/stratus/sbin/certificate_installer.sh -c/tmp/cert.pem -
   k/tmp/key.pem
   ```

To revert to the default certificate

```
$ sudo /opt/stratus/sbin/certificate_installer.sh -r
```

To display information about the currently used certificate

```
$ sudo /opt/stratus/sbin/certificate_installer.sh -d
```

7

Chapter 7: Using the ztC Endurance Console

The ztC Endurance console is a browser-based interface that provides management and monitoring of a ztC Endurance system from a remote management computer. For an overview of the console, see "ztC Endurance Console" on page 218.

For information on pages within the ztC Endurance console, see the following topics:

- "Dashboard Page" on page 224
- "Configuring the Mail Server" on page 227
- "Hardware Page" on page 229
- "Volumes Page" on page 238
- "Specifying Owner Information" on page 238
- "Configuring Users" on page 244
- "Alert History Page" on page 246
- "Support Logs Page" on page 247
- "Configuring e-Alerts" on page 248
- "Configuring SNMP Settings" on page 250
- "Configuring Remote Support Settings" on page 264
- "Configuring Internet Proxy Settings" on page 267

ztC Endurance Console

The ztC Endurance console is a browser-based interface that provides management and monitoring of a ztC Endurance system from a remote management computer. You can perform many administrative operations from the console because it provides access to the system as a whole as well as to individual modules, components of modules, and other resources.

For information on the requirements of the remote management computer that runs the ztC Endurance console, see " ztC Endurance Console Requirements" on page 219.

Using the options listed in the left pane of the ztC Endurance console, you can perform a variety of administrative functions:

- Read system alerts from the Dashboard. See "Dashboard Page" on page 224.

- Configure the mail server. See "Configuring the Mail Server" on page 227.

- View information about the system and its components. See "Hardware Page" on page 229.

- Monitor resources for volumes, including their state, name, data synchronization status, and size. See "Volumes Page" on page 238.

- Specify the name and contact information for an administrator or owner of the system. See "Specifying Owner Information" on page 238.

- Register your system. See "Registering the System" on page 239.

- Configure users. See "Configuring Users" on page 244.

- Change passwords. See "Changing Passwords on a ztC Endurance System" on page 228.

- View alerts. See "Alert History Page" on page 246.

- View support logs. See "Support Logs Page" on page 247.

- Configuring notifications. See "Configuring e-Alerts" on page 248 and "Configuring SNMP Settings" on page 250.

- Enable OPC functionality. See "Configuring OPC Settings" on page 257.

- Configure remote support settings to send alerts to Penguin Solutions Customer Service via the Active Service Network (ASN) when an event requires attention. See "Configuring Remote Support Settings" on page 264.

- Configure proxy settings to support notification messaging and remote support access features. See "Configuring Internet Proxy Settings" on page 267.

- Upload software upgrade kits. See "Upgrade Kits Page" on page 268.

In addition, you can use the ztC Endurance console to:

- Edit your user information. See "Editing Your User Information" on page 223.

- View your system's serial number by hovering the cursor over the ztC Endurance model number at the top of the left pane.

- View the management IP addresses, the system's software version number, and asset ID in the banner at the top of the ztC Endurance console.

Related Topics

- "Logging In to the ztC Endurance Console for the First Time" on page 220

- "Logging In to the ztC Endurance Console" on page 221

- "Using the ztC Endurance Console" on page 217

ztC Endurance Console Requirements

The ztC Endurance console provides browser-based remote management of the ztC Endurance system, including its processing, storage, networks, and power.

- Your computer must be able to access the subnet containing the ztC Endurance management network, which is enabled on the port labeled **1G MGMT**. See Figure 1-4 for the location of this port.

- Use a supported browser. See "Compatible Internet Browsers" on page 219.

Related Topics

- "ztC Endurance Console" on page 218

- "Using the ztC Endurance Console" on page 217

Compatible Internet Browsers

A browser is used to connect to the ztC Endurance console. Use only browsers that are compatible with ztC Endurance systems. Using an incompatible browser can result in some rendering problems and the omission of some wizards. You can view the ztC Endurance console in dark mode by adjusting the settings in your browser window, but the settings differ depending on the browser. See "To enable dark mode in the ztC Endurance console" on page 220 for more information.

Table 7-1 lists browsers that are compatible with ztC Endurance systems.

Table 7-1: Compatible Browsers

Compatible Browsers	Release
Microsoft Edge	118 or later
Mozilla® Firefox®	118 or later
Google® Chrome™	118 or later

To enable dark mode in the ztC Endurance console

In Microsoft Edge and Firefox, browser settings automatically apply to the ztC Endurance console, so if you select dark mode in **Settings**, the console appears in dark mode.

In Chrome, the browser's appearance settings do not automatically apply to the ztC Endurance console. To activate dark mode in the console while using Chrome, perform the following procedure:

1. With the ztC Endurance console open, click the three vertical dots menu at the top of the browser page. Select **More tools**, then click **Developer tools**.

2. In **Developer tools**, click the smaller vertical dots menu and select **More tools**. Click **Rendering**.

3. In the **Rendering** tab at the bottom of the Developer tools window, scroll down to locate the **Emulate CSS media feature prefers-color-scheme** option. Select **prefers-color-scheme: dark**.

Logging In to the ztC Endurance Console for the First Time

> **Note**: If this is **not** your first time logging in to the ztC Endurance console, see "Logging In to the ztC Endurance Console" on page 221 instead.

When deploying the system, log in to the ztC Endurance console to change the default password.

To log in to the ztC Endurance console for the first time

1. From a networked PC or laptop, type the secure version (**https**) of one of the two management IP addresses of your ztC Endurance system into a browser address bar.

> **Note**: If a security message appears, proceed to the web site. You can add a security exception later, to allow the site to load without the message "Configuring Secure Connections."

2. The login page of the ztC Endurance console appears. Enter **admin** for the **Username** and **admin** for the **Password** (or other credentials, if provided), and then click **LOGIN**.

3. When you see the **Warning** message notifying you to change your password, click **OK**.

> **Notice**: You must change the password for **admin** now, for security purposes. You can change it again later when you update the default user profile for the **admin** account. You make these changes on the **Users** page (see "Configuring Users" on page 244).

4. On the **Change Password** window, provide the requested password information and click **Continue**. The password must conform to the system's password policy (for information, see Password Policy).

5. When you see a message saying that your password reset was successful, click **Finish**.

6. When the login page appears, enter **admin** for the **Username** and your new password, and then click **LOGIN**.

The ztC Endurance console appears and the initial login is complete. Bookmark or otherwise make note of the system IP address for use when logging in to the console in the future.

Perform additional tasks in "Post-Installation Tasks" on page 183, if necessary.

Related Topics

- *ztC Endurance 3xx0, 5xx0, 7xx0, and 9xx0 Systems: Deploying a System* (R002C)

- "Logging In to the ztC Endurance Console" on page 221

- "ztC Endurance Console" on page 218

Logging In to the ztC Endurance Console

> **Note**: If this is your first time logging in to the ztC Endurance console, see "Logging In to the ztC Endurance Console for the First Time" on page 220 instead.

Log in to the ztC Endurance console to manage the ztC Endurance system. Using the console, you can manage the system, including its processing, storage, networks, and power. You can also view alerts and perform other administrative tasks.

Notes:

1. A login session times out after 10 minutes, if unused.

2. The system has a limit of 10 login sessions.

3. A username cannot start with a special character and is limited to 64 characters.

4. Passwords must conform to the system's "Password Policy" on page 223.

To log in to the ztC Endurance console

1. Type one of the two management IP addresses of your ztC Endurance system or the fully qualified domain name (FQDN) into a browser address bar:

   https://*IP_address*

   OR

   https://*FQDN_name*

2. When the login page appears, enter your **Username** and **Password**.

   If you have forgotten your password, click **Forgot Password?** and the **Reset Password** page appears. Enter the requested information to reset your password.

   Note: Resetting a password requires that you have an email account on the system, with an email address, as configured in your user account (see "Managing User Accounts" on page 245). If you are unable to receive email, contact your system administrator.

To reset your password

> **Note**: To receive email when resetting your password, the mail server must be configured. See "Configuring the Mail Server" on page 227.

   a. When the **Reset Password** page appears, enter your **Username** and click **Continue**. An email is sent to the email address listed with your local user account. The email contains a link to a reset password page.

   b. In your email account, open the email with the reset-password link, and click the link. The **Reset Password** page reappears.

   c. For **New Password** and **Confirm Password**, type your new password. The new password must conform to the system's "Password Policy" on page 223.

      Click **Continue**.

   d. A page appears, with a message that the reset was successful and that you can log in to the system with your new password. Click **Finish**.

3. Click **LOGIN**.

Password Policy

The system's password policy requires that your password meet these conditions:

- Its minimum length is 8 characters, and its maximum length is 72 bytes.

- It must contain both upper- and lowercase characters.

Related Topics

- "Logging In to the ztC Endurance Console for the First Time" on page 220

- "ztC Endurance Console" on page 218

Editing Your User Information

Edit your user information (that is, your user profile) by changing your user name, email address, first name, last name, or password.

To edit your user information

1. Click your user name in the upper right-hand corner of the console.

   The **Edit User** dialog box opens.

2. Enter or modify values for the following:

- **User Name**

- **Email Address**

- **First Name** (optional)

- **Last Name** (optional)

- **Password**

> **Note**: Passwords must confirm to the system's "Password Policy" on page 223.

- **Confirm Password**

3. Click **Save**. (Or click **Cancel** to cancel the changes.)

Related Topics

- "ztC Endurance Console" on page 218

Dashboard Page

The **Dashboard** page:

- Displays a summary of outstanding alerts on the ztC Endurance system.

- Allows you to perform a total system shutdown by clicking the **Shutdown** button.

> **Notice**: Do not click the **Shutdown** button unless you are familiar with the topics discussed in "System Power" on page 86.

The **Dashboard** page opens when you log in to the ztC Endurance console. To open it from other pages, in the left-hand navigation panel, click **Dashboard** under **System**.

On the **Dashboard** page:

- The masthead at the top of the page displays registration status information, system-alert status, the name of the current user, and the **Logout** option.

- Below the masthead, the ztC Endurance model number appears. To view your system's serial number, hover the cursor over the ztC Endurance model number. The system hostname, the management IP address of compute module A, the management address of compute module B, the software version number (RH refers to Red Hat), and the asset ID also appear below the masthead, along with the **Shutdown** button and a **Help** button that takes you to the ztC Endurance Help pages.

- The bottom pane displays a list of alerts associated with the components shown in the system diagram.

- The left-hand navigation panel displays the pages that allow you to configure or view information about the system.

To display additional information about outstanding alerts, click an alert symbol (for example, ⓘ) in the system diagram, or click an entry in the list of alerts below the system diagram. Alert lists may appear in tabs such as **All**, **System**, and **Ignored**, as well as in tabs for each component shown in the system diagram. These tabs may appear below the system diagram, depending on the alerts. The alert information includes:

- The component associated with the issue (for example, the ztC Endurance system or a module).

- A description of the activity or task that requires attention.

- The ability to ignore the issue (only if **Ignore** appears in the **Action** column). You can restore the alert by clicking the **Ignored** tab and then clicking **Restore** in the **Action** column for the alert.

- The reason the issue should be resolved by clicking the row with the alert. The information will appear in the **Information** column.

Resolve active alerts as soon as possible (see "Resolving Outstanding Alerts on the Dashboard" on page 226).

Understanding the ztC Endurance System Diagram

The system diagram displays a graphical representation of system status. A star symbol (⭐) indicates the active module.

The system diagram also shows which disk-drive slots and PCIe-card slots are populated. If a disk-drive slot is grayed-out in the diagram, that slot is empty. Similarly, if the diagram shows a PCIe-card slot with slot covers instead of a PCIe card with ports, that slot is empty.

Alert symbols, if present, represent informational or critical alerts that require attention. Click an alert symbol to display information about the alert.

Alert Symbols

General alert symbols, in descending order of severity, are:

- Critical or Serious (⛔)

- Moderate (🛡)

- Minor (⚠)

- Informational ( )

- Good ( )

Device-specific alert symbols include:

- PSU symbols:

    - : Reports an alert on a failing or degraded PSU

    - : Reports a warning alert on a PSU (for example, if the PSU is simplex)

- Disk symbols:

    - : Reports an alert on a failing or degraded physical disk

    - : Reports a warning alert on a physical disk (for example, if the disk is simplex)

- Network symbols:

    - : Reports an alert on a failing or degraded teamed network port.

    - : Reports a warning alert on a teamed network port (for example, if the network port is simplex)

Related Topics

- "ztC Endurance Console" on page 218

Resolving Outstanding Alerts on the Dashboard

Resolve any outstanding alerts that appear on the **Dashboard** page.

To resolve outstanding alerts

If an alert appears in the ztC Endurance system diagram on the ztC Endurance console **Dashboard** page, click the alert. Information about it appears in the bottom pane. Your options are as follows:

- Resolve the alert.

    For instance, if you see the message **Support Notification service should be enabled to ensure the best possible support from Stratus**, enable the support notification service. See "Configuring Remote Support Settings" on page 264.

- For minor alerts, click **Ignore** (beneath the **Action** column) to ignore the alert and remove it from the list. Clicking **Ignore** hides the alert. You cannot ignore more serious alerts.

    To restore the ignored alert to the list, click the **Ignored** tab above the alerts list, and then click **Restore**, under the **Action** column.

Related Topics

Configuring the Mail Server

Configure the mail server to enable the ztC Endurance system to send email when, for example, a user needs to reset a password, generate a test alert, or receive alerts via email (e-Alerts).

To configure the mail server

> **Note**: If you change any mail-server settings, you **must** re-enter the mail-server password if authentication is enabled.

1. In the left-hand navigation panel, click **Mail Server Configuration** under **System**.

2. On the **Mail Server Configuration** page, click the **Enable Mail Server** box. Boxes for specifying or selecting the following settings appear:

   - **SMTP Server** (required)–Enter the name of the Simple Mail Transfer Protocol (SMTP) server that your company uses to send email.

   - **Port Number**–Enter the port number to use when sending e-Alerts. If no port number is specified, the default SMTP port 25 is used.

   - **Sender's Email Address**–Enable e-Alert delivery by specifying a valid sender's email address in either of the following cases:

     ○ You have not specified a DNS server on the ztC Endurance system **and** your SMTP server is not configured to accept domain literals (`From` addresses in the form `noreply@IP_address`).

     ○ You want the e-Alert to provide a different sender's email address (for example, `noreply@company.com`).

   Any email address that the SMTP server accepts is sufficient.

   - **Encrypted Connection**–Select a value from the drop-down menu for the encryption protocol that the SMTP server requires:

     > **Note**: Regardless of whether you select **TLS** or **SSL**, the ztC Endurance software selects whichever protocol is best at the time of the connection.

- ○ **None** for no encryption. By default, port number 25 is used.

- ○ **TLS** for the Transport Layer Security (TLS) protocol. For TLS, Stratus recommends that you specify 587 for **Port Number**, though 25 is used by default. However, check with your network administrator to make sure that 587 is the correct port number.

- ○ **SSL** for the Secure Sockets Layer (SSL) protocol. For SSL, Stratus recommends that you specify 465 for **Port Number**, though 25 is used by default. However, check with your network administrator to make sure that 465 is the correct port number.

- ■ **Enable Authentication**—Click this box if the SMTP server requires authentication to send email. Then, type the **Username** and **Password** for the SMTP account.

> **Note**: If authentication is enabled (because the **Enable Authentication** box is already checked or because you have just checked it) and you change any mail server settings, you **must** re-enter the mail-server password.

3. Click **Save** (or click **Cancel** to restore the previously-saved values).

Related Topics

- "ztC Endurance Console" on page 218

Changing Passwords on a ztC Endurance System

After installation, use the **Change Passwords** page on the ztC Endurance console to change the following user-configurable passwords:

- **BMCs - Admin** – BMC Web console administrator account password.

- **BMCs - Sysadmin** – Password for the Debug console `sysadmin` user account. This is automatically set by the ztC Endurance system during initial installation, but you can change it. Most users never need to access the BMC Debug console.

- **Redfish** – Password for the BMC Redfish administrator account. This is automatically set by the ztC Endurance system during initial installation, but you can change it. Most users never need to access the Redfish administrator account.

- **Management User** – `zenadmin` account on the Standby OS.

See "Gathering Information" on page 146 to learn more about the BMC Web console administrator account password and the `zenadmin` password.

> **Note**: Changing a password requires placing the standby compute module in maintenance mode, and changing the Management User password requires a Smart Exchange. The system prompts you to confirm these actions before proceeding.

Use the ztC Endurance console to change the passwords for accounts on the ztC Endurance system.

To change passwords

1. Log in to the ztC Endurance console as described in "Logging In to the ztC Endurance Console" on page 221.

2. Under **Administrative Tools**, click **Change Passwords**.

3. On the **Change Passwords** page, select the radio button for the password you want to change: **BMCs - Admin**, **BMCs - Sysadmin**, **Redfish**, or **Management User**.

4. If the **Old Password** field is present, enter the existing password.

5. In the **New Password** field, enter the new password.

6. In the **Confirm New Password** field, enter the new password again.

7. Click **Save** to save your changes, or click **Cancel** to cancel the change-password operation.

Hardware Page

The **Hardware** page displays information about the hardware components of a ztC Endurance system. The top pane lists the state, activity type, name, and serial number of each compute module, I/O module, storage module, and PSU in the system. The bottom pane displays detailed information about the top-pane component that is currently selected.

When you select tabs in the bottom pane, the information displayed is generally self-explanatory. However, additional information about the Memory, Network, Fibre Channel, and Storage tabs is provided later in this section.

In addition to viewing information about hardware components, you can use the **Hardware** page to do the following:

- Identify a hardware component by clicking a button to flash its ID LED (see "Identifying a System or Component" on page 234).

- Place a compute module into maintenance mode, shut it down, and then return it to service (see "Using Maintenance Mode" on page 235).

- Go to the BMC login page. From the compute module's **BMC** tab, click an IP address to go to the BMC login page (see "Logging In to the BMC Web Console" on page 273).

- View operating-system information (name, version, IP address, and hostname) by clicking the compute module's **OS** tab.

Hardware States

The **State** values show the severity of the component's state. Values are:

- Informational ( ): Hardware slot not actively involved in the system.

- Normal, in sync, or OK state ( ): Operating normally.

- Minor, warning, or inconsistent state ( ): User attention recommended.

- Moderate state ( ): Module is simplex; user attention required.

- Broken, failed, or severe state ( ): Component not operational; user attention required. **Activity** text may provide additional information about this state.

- Syncing ( ): Component is working to become redundant again or the compute module is performing a Smart Exchange.

- Maintenance mode ( ): Component is in maintenance mode. See "Using Maintenance Mode" on page 235.

Hardware Activity Types

The **Activity** values show the component's current activity type. Values are:

- **Running** ( ): Component provides an operating environment and is running normally. This value applies only to the active or standby compute module.

- **Connected** ( ): Component is actively connected to the active compute module and is operating normally. This value applies to storage modules, I/O modules, and PSUs.

- **Shutting Down** ( ): Compute module is shutting down. See "Shutting Down and Restarting a Compute Module" on page 237.

- **Shutdown** ( ): Compute module has shut down.

- **Smart Exchange** ( ): Active compute module is performing a Smart Exchange to the other compute module.

- **Recovering** ( ): Compute module is transitioning.

- **Booting** ( ): Compute module is booting.

- **Failed** ( ): Component has failed and is not operational.

- **Disconnected** (✖): Component is not connected to the active compute module.

- **Off** (⏻): Component is powered off.

To learn more about the Memory tab

On the **Hardware** page, you can learn more about a specific compute module's memory by selecting the compute module and then clicking its **Memory** tab in the lower pane. The **Memory** columns have the following meanings:

- **State**: Displays the state of the RDIMM. For example, a green checkmark in the **State** column indicates that the RDIMM's state is good. See "Hardware States" on page 230 for more information about the hardware states and their associated icons.

- **Name**: Displays the name of the RDIMM (for example, **DIMM8**).

- **Slot**: Displays which of the two RDIMM slots on a channel an RDIMM is physically plugged in to. Because ztC Endurance systems implement only one slot per channel, you can ignore the values in this column.

- **Socket**: Displays the processor to which an RDIMM is connected. For example, a value of **1** indicates that the RDIMM is connected to processor 1.

- **Channel**: Indicates the number of the channel (that is, a physical set of wires) between the processor and the RDIMMs. Each processor in a ztC Endurance system has 8 RDIMM channels.

- **Capacity**: Displays the size of the RDIMM, in gigabytes (for example, **32 GB**).

See "Memory in ztC Endurance Systems" on page 22 for more information about RDIMMs.

To learn more about the Network tab

On the **Hardware** page, you can learn more about a system's Ethernet PCIe adapters.

- To learn about a specific compute module's embedded Ethernet PCIe adapters, select the compute module and then click its **Network** tab in the lower pane.

- To learn about a specific I/O module's Ethernet PCIe adapters, select the I/O module and then click its **Network** tab in the lower pane.

The **Network** columns have the following meanings:

- **State**: Displays the state of the component. For example, a green checkmark in the **State** column indicates that the PCIe adapter's state is good. See "Hardware States" on page 230 for more information about the hardware states and their associated icons.

- **Location**: Displays the location of the component. For example, in a compute module, **embedded-A** means that the embedded PCIe adapter is located in PCIe slot A. In an I/O module, **slot B/2** means that the PCIe adapter is located in I/O module B, slot 2.

- **Name**: Displays the name of the component (for example, **ethA0601**).

- **Port**: Displays which port the component is connected to (for example, **1**).

- **PCI ID**: Indicates the component's PCIe address identifier on the system (for example, **0000:ce:00.1**).

- **Speed**: Indicates the speed of the component (for example, **10 Gbps**).

- **MAC Address**: Displays the component's MAC address (for example, **34:ad:61:fb:19:d5**).

- **Bonds**: If the component is a member of a bond, displays the bond name (for example, **bond0**). If the component is not a member of a bond, the field appears blank.

To learn more about the Fibre Channel tab

On the **Hardware** page, you can learn more about a system's Fibre Channel (FC) PCIe adapters by selecting the I/O module and then clicking its **Fibre Channel** tab in the lower pane.

> **Note**: The **Fibre Channel** tab appears only in systems that contain an FC PCIe adapter.

The **Fibre Channel** columns have the following meanings:

- **State**: Displays the state of the component. For example, a green checkmark in the **State** column indicates that the FC PCIe adapter's state is good. See "Hardware States" on page 230 for more information about the hardware states and their associated icons.

- **Location**: Displays the location of the component. For example, **slot A/2** means that the FC PCIe adapter is located in I/O module A, slot 2.

- **Name**: Displays the name of the component (for example, **host3**).

- **PCI ID**: Indicates the component's PCIe address identifier on the system (for example, **0000:cc:00.0**).

- **Speed**: Indicates the speed of the component (for example, **32 Gbps**).

- **WW Node Name**: Displays the unique identifier for the component (for example, **20:00:f4:c7:aa:3c:10:ad**).

- **WW Port Name**: Displays the unique identifier for the port on the component (for example, **21:00:f4:c7:aa:3c:10:ad**).

To learn more about the Storage tab

On the **Hardware** page, you can learn more about a specific storage module's disks by selecting the storage module and then clicking its **Storage** tab in the lower pane. The **Storage** columns have the following meanings:

- **State**: Displays the state of the disk. For example, a green checkmark in the **State** column indicates that the disk's state is good. See "Hardware States" on page 230 for more information about the hardware states and their associated icons.

- **Location**: Displays the location of the disk. For example, **slot A/6** indicates that the disk is located in storage module A, slot 6.

- **Name**: Displays the name of the disk (for example, **nvme0n1**). If the disk slot is empty, **Empty** is displayed.

- **PCI ID**: Indicates the disk's PCIe address identifier on the system (for example, **0001:74:00.0**).

- **Availability**: Indicates if the disk is redundant (that is, paired with another physical disk). If the disk slot is empty, **Unused** is displayed.

- **Logical ID**: Displays the name of the RAID array volume if the disk is a member of an array (for example, **md20**). If the disk slot is empty, **Unconfigured** is displayed.

Related Topics

- "Removing and Replacing System Components" on page 101

- "Dashboard Page" on page 224

- "Disk Drives" on page 27

- "Memory in ztC Endurance Systems" on page 22

- "PCIe Adapters" on page 30

- "System CRU Modules" on page 20

- "General Servicing Information" on page 101

- "ztC Endurance Console" on page 218

Identifying a System or Component

You can identify:

- A specific ztC Endurance system in a cabinet

- A specific component within a ztC Endurance system

To identify a ztC Endurance system

Press the **ID** button on the compute module (see Figure 1-5). The blue ID LED on the control-panel ear (see Figure 4-2) illuminates for 30 seconds.

To identify a component in a ztC Endurance system

1. In the left-hand navigation panel, click **Hardware** under **System**.

2. In the top pane of the **Hardware** page, select the component that you want to identify and then

    click the **Identify** ( ) button. The blue ID LED for that component illuminates for 30

    seconds. See "Status LEDs" on page 90 for the location of each component's ID LED.

Related Topics

- "Hardware Page" on page 229

- "Dashboard Page" on page 224

- "Compute Module" on page 21

Using Maintenance Mode

When you place a compute module into maintenance mode, it goes offline for service. While a compute module is in maintenance mode, operations such as monitoring, recovery, automatic power-on, and so on cannot occur on that compute module.

If you place the **active** compute module into maintenance mode, the system automatically performs a Smart Exchange.

What is Smart Exchange?

Smart Exchange is a Stratus proprietary process that maintains system availability and protects against data loss. If the ztC Endurance system detects a degradation or failure within the active compute module, Smart Exchange:

- Automatically activates the standby compute module and copies over in-process data

- Reroutes all I/O and storage connected to the original active compute module over to the newly activated compute module

Smart Exchange occurs while the system is running and except for a very brief pause (usually, less than one second), is undetectable to the operating system and customer applications.

You must manually remove the compute module from maintenance mode to return it to service.

Note that when the active compute module enters maintenance mode, the standby compute module becomes active after the Smart Exchange completes. During this operation, communications between the active and standby compute modules are disabled, and the operating system and applications running on the original active compute module are migrated to the new active compute module. This protects the operating system and running applications from any potential disruption caused by the operation.

> **Notices**:
>
> 1. A compute module must be in maintenance mode before you shut it down. For information about shutting down a compute module, see "Shutting Down and Restarting a Compute Module" on page 237.
>
> 2. The system is not fault tolerant when a compute module is in maintenance mode. For continuous uptime, finalize service as soon as possible so that the compute module can exit maintenance mode and go back online.

To place a compute module into maintenance mode

1. From the **Hardware** page, select a compute module.

2. Click **Work On**.

3. Click **Yes** on the **Confirm** window.

When the compute module is in maintenance mode, its state appears as **Maintenance Mode**, and its **Activity** displays as **Running**. At this point, it is safe to shut down the compute module if necessary.

To remove a compute module from maintenance mode

1. From the **Hardware** page, select the compute module that is in maintenance mode.

2. Click **Finalize**, which removes the compute module from maintenance mode.

   When the compute module exits maintenance mode, its state appears as **Normal**, and its **Activity** displays as **Running**.

Related Topics

- "ztC Endurance Console" on page 218

- "Hardware Page" on page 229

- "Shutting Down and Restarting a Compute Module" on page 237

- "Removing and Replacing a Compute Module" on page 106

Shutting Down and Restarting a Compute Module

From the **Hardware** page, you can shut down and restart a compute module. You do this to perform service operations that require physically removing a compute module from the system.

> **Notice**: A compute module must be in maintenance mode before you shut it down. For information about entering a compute module into maintenance mode, see "Using Maintenance Mode" on page 235.

To shut down a compute module

1. From the **Hardware** page, select a compute module that is in maintenance mode.

2. Click **Shutdown**.

3. Click **Yes** on the **Confirm** window.

   When the compute module has been shut down, its state appears as **Maintenance Mode**, and its **Activity** displays as **Shutdown**. At this point, it is safe to remove the compute module.

To restart a compute module

1. From the **Hardware** page, select a stopped compute module.

2. Click **Boot**.

3. Click **Yes** on the **Confirm** window.

   After the compute module completes its restart, its state still appears as **Maintenance Mode**, and its **Activity** displays as **Running**.

4. Remove the compute module from maintenance mode.

Related Topics

- "ztC Endurance Console" on page 218

- "Hardware Page" on page 229

- "Using Maintenance Mode" on page 235

- "Removing and Replacing a Compute Module" on page 106

Volumes Page

The **Volumes** page displays information about logical storage volumes created from the physical devices in the storage modules. To open this page, in the left-hand navigation panel, click **Volumes** under **System**.

The **Volumes** page displays the following columns with information about volumes in the top pane:

- State

- Name

- Disk Synchronization. If a volume is synchronizing, this column displays the percentage complete and the estimated time of completion.

- Size

- Mounts

Click the name of a volume in the top pane of the **Volumes** page to display additional information about the volume in the bottom pane. This information includes the physical members of the volume as well as their states.

Related Topics

- "Disk Drives" on page 27

- "ztC Endurance Console" on page 218

Specifying Owner Information

Specify the name and contact information for an administrator or owner of the ztC Endurance system to make this information available for support purposes. This contact information is included in any email sent by e-Alerts.

This contact information is available in the ztC Endurance console and provided in response to Simple Network Management Protocol (SNMP) requests.

> **Note**: The information you specify during Step 1 and Step 2 of the registration portal procedure (see "Registering the System" on page 239) is automatically supplied as the system owner information on the **Owner Information** page. However, if you have already specified system owner information on the **Owner Information** page and then register your system, the information specified during registration does not overwrite the system owner information.

To specify system owner information

1. In the left-hand navigation panel, click **Owner Information** under **System**.

2. On the **Owner Information** page, supply information in the **Full Name**, **Phone Number**, **Email Address**, and **Site Address** fields.

3. Click **Save**.

Related Topics

- "ztC Endurance Console" on page 218

Registering the System

Registering a ztC Endurance system:

- Helps Penguin Solutions Customer Service deliver responsive, expedited service, should you contact Penguin Solutions with a potential issue or need a replacement part.

- Registers you for the **Penguin Solutions Customer Service Portal** at https://service.stratus.com, giving you access to technical resources–including Knowledge Base articles, Customer Service publications, and software downloads–that are available only to registered users.

To use these services, it is important to register the system soon after you log in to the ztC Endurance console for the first time. However, you can register the system later. You can do so on a system with or without Internet access.

> **Note**: The information you specify during Step 1 and Step 2 of the registration portal procedure is automatically supplied as the system owner information on the **Owner Information** page (see "Specifying Owner Information" on page 238). However, if you have already specified system owner information on the **Owner Information** page and then register your system, the information specified during registration does not overwrite the system owner information.

On a system that does not have Internet access, you need to move a file between the location of the ztC Endurance console (which does not have Internet access) and a location with Internet access. Two possible methods for moving the file follow.

- A USB flash drive—You move a USB flash drive between a management PC (which can connect to the system) and a computer with Internet access.

- A mobile device such as a laptop or smart phone—You move a mobile device between a location where you can log in to the ztC Endurance console and a location with Internet access.

> **Prerequisite**: Before you register the system, read To complete the registration portal steps to ensure that you have all required information.

To complete the registration portal steps

Step 1: General Information–Enter the following information:

- **First Name** and **Last Name**

- **Company Email**–Provide the email address of the company that is the final deployment site. Do not provide a personal email address.

You also need to review and accept the **Service Terms**.

Step 2: Location Info–Enter the following information:

- **End User Company Name**–Provide the full name of the company where the system will be deployed.

- **Deployment Shipping Address**–Provide the complete address for replacement-part shipments. Use the address of the company that is the final deployment site. Do not provide a post-office box. Fields are:

  - **Address 1** and **Address 2**

  - **City**, **State**, **Postal Code**, and **Country**

  - **Special Instructions** (for example, "always deliver to loading dock 2")

Step 3: Contact Details–Enter the following information:

- **Primary Technical Contact** and **Secondary Technical Contact**–Provide the names of the technical contacts who will be communicating with your authorized Penguin Solutions service representative.

- **Service Renewal Contact**–Provide the name of the person who is responsible for handling annual service-agreement renewals.

For each contact, enter **First Name**, **Last Name**, **Email Address**, **Desk Phone**, and **Mobile (optional)**. You can add more contacts later using the **Penguin Solutions Customer Service Portal** at https://service.stratus.com.

After you click **Next** at the bottom of the page, Stratus verifies the information.

If there is a problem with the information, a **Problem Encountered** pane appears, which describes the problem. Click **Back** to fix the problem, if possible. If a problem still exists, click **Next** to continue, allowing you to download a file that enables you to complete registration. To help resolve the problem and ensure that your account is set up properly, your authorized Penguin Solutions service representative will contact you.

An **Information Verification** page appears, allowing you to review the information. Click **Back** to change any information. Click **Next** to submit the information and complete registration.

Step 4: License Key–For a system with Internet access, check the **License** page of the ztC Endurance console to confirm that the license has been installed. For a system without Internet access, click **Download License** to download the license-key file, which you will install on the ztC Endurance system. Make note of the location where you download the file.

To register a system and acquire a license

On a system with Internet access

1. In the ztC Endurance console, do one of the following:

   a. Click **License**, under **System**, in the left-hand navigation panel.

   b. Click **Register for License** in the masthead.

2. For **Online License Registration and Activation**, click **Register Online** to open a new browser tab with the Stratus registration web portal. Then, complete the [registration web portal steps](#).

   At **Step 4**, check the **License** page of the ztC Endurance console to confirm that the license has been installed. If needed, click **Check License Now** next to **Online License Check**. When **STATUS** changes to **License installed**, the license has updated successfully, and registration is complete.

   If the license does not update successfully within 5 minutes, click **Download License** in the registration web portal to download the license-key file. Continue with the next step.

3. On the **License** page of the ztC Endurance console, under **Offline License Check and Manual License Installation**, click **Choose File** under **Install an Activated License Key to the System**. Navigate to the location where you saved the file, select the file, and click **Open**.

4. On the **License** page of the ztC Endurance console, click **Upload** to upload the file to the system. If needed, click **Check License Now** next to **Online License Check**. When **STATUS** changes to **License installed**, the license has updated successfully, and registration is complete.

On a system without Internet access

If a system does not have Internet access, you need to move a file between the location of the ztC Endurance console (which does not have Internet access) and a location with Internet access. The following procedure describes one method, though other methods are possible.

On a computer or mobile device with access to the ztC Endurance console

1. If using a management PC, insert a USB flash drive into a USB port.

    If using a mobile device, ensure that it has access to the ztC Endurance console.

2. In the ztC Endurance console, click **License**, under **System**, in the left-hand navigation panel.

3. For Step 1, **Offline License Registration via URL File** (beneath the **Offline License Registration and Manual License Installation** bar), click **Download URL File** and save the *register_site_file*.html file to the USB flash drive or mobile device. If using a USB flash drive, remove it.

4. Go to a location with Internet access.

In a location with Internet access

1. If using a USB flash drive, insert it into a USB port of the computer with Internet access.

2. Navigate to the file you saved, and click the file name. A browser opens the file and is redirected to the Stratus registration web portal. Complete the [registration web portal steps](#).

    At **Step 4**, download the permanent license-key file and save it to the USB flash drive or mobile device. If using a USB flash drive, remove it.

3. Return to the location with access to the console.

On a computer or mobile device with access to the ztC Endurance console

1. If using a USB flash drive, insert it into a USB port on the management PC.

    If using a mobile device, ensure that it has access to the ztC Endurance console.

2. In the ztC Endurance console, click **License**, under **System**, in the left-hand navigation panel.

3. For Step 2, **Install an Activated License Key to the System** (beneath the **Offline License Registration and Manual License Installation** bar), click **Choose File**. Navigate to the location where you saved the license-key file, select the file, and click **Open**.

4. Click **Upload** to upload the file to the system.

Related Topics

Configuring Users

Use the **Users** page to add, modify, or remove user accounts on the ztC Endurance system. You can select a user and view the time when the user's password was last updated. An administrator can also use the page to force a selected user to change the user's password on the next login.

To open this page, in the left-hand navigation panel, click **Users** under **Administrative Tools**.

To manage user accounts

To add a new user, click **Add** in the lower pane. On the **Add a User** window, supply the requested information and click **Save**.

To modify an existing user, click the name of a user account and click **Edit**. On the **Edit User** window, modify the information and click **Save**.

To remove an existing user, click the name of a user account and click **Remove**. On the **Confirm** window, click **Yes**.

To view the time when a user last changed the user's password, look at the **Last Password Update Time** column for a selected user. To force a user to change the user's password on the next login, an administrator selects the user and then clicks **Expire Password**.

For more information, see .

To sort and locate user accounts

If you have a large number of accounts, you can click a column heading to sort the accounts by parameter. You can sort accounts by **Type**, **Username**, **First Name**, **Last Name**, **Email**, **Role**, or **Last Password Update Time**.

Related Topics

-

Managing User Accounts

You add, edit, or remove users, specify passwords, and assign user roles to user accounts on the **Users** page in the ztC Endurance console. You can also select a user and view the time when the user's password was last updated, and an administrator can force a selected user to change the user's password on the next login.

User roles are:

- **Administrator**: full system administrator privileges

- **Platform Manager**: system administrator privileges except for adding, removing, and modifying users

- **Read-only**: ability to view but not to change system configuration or to install system software

For the procedures below, begin by opening the **Users** page: in the left-hand navigation panel, select **Users** under **Administrative Tools**.

To add a user account

1. In the lower pane, click **Add**.

2. In the **Role** drop-down window, select **Administrator**, **Platform Manager**, or **Read-only**.

3. Enter values for the **User Name**, **Email Address**, **First Name**, **Last Name**, **Password**, and **Confirm Password** fields. User names may be from 1 to 64 characters long, and must include no white space. Passwords must conform to the "Password Policy" on page 223 of the system. Click **Save**.

To edit a user account

1. Select the account you want to edit.

2. In the lower pane, click **Edit**.

3. Change the user's information, as necessary. For example, to change a user's role, in the **Role** drop-down window, select **Administrator**, **Platform Manager**, or **Read-only**.

4. Click **Save**.

To force a user to change the user's password

1. Select the user whose password you want to expire.

2. Click **Expire Password**.

3. Click **Yes** in the **Confirm** dialog box.

To remove a user account

1. Select the account to remove.

2. Click **Remove** in the lower pane.

3. Click **Yes** in the **Confirm** dialog box.

> **Notes**:
>
> 1. You cannot delete the default **admin** account if it is the only administrator account, although you should change its name and password by editing the account.
>
> 2. You must specify an email address for each user account, including **admin**, to enable the forgot-password feature. If a user account does not include an email address, and the user clicks the **Forgot Password?** link on the console login page, the system sends an email to *user@example*.com.

Related Topics

- "ztC Endurance Console" on page 218

- "Configuring Users" on page 244

Alert History Page

The **Alert History** page displays messages about events on the ztC Endurance system. To open this page, in the left-hand navigation panel, click **Alert History** under **Alerts & Logs**.

> **Note**: Support notification alerts, e-Alerts, and SNMP traps are generated only when you enable them in the ztC Endurance console. For information, see:
>
> - "Configuring Remote Support Settings" on page 264
>
> - "Configuring e-Alerts" on page 248
>
> - "Configuring SNMP Settings" on page 250

To view alert information, scroll through the alerts, which are, by default, listed in reverse chronological order. On the upper pane, click an alert to display the following information in the lower pane:

- The date and time the alert occurred

- Information about the problem and resolution (if available)

- Whether a **Support Notification**, **e-Alert**, or **SNMP Trap** was sent for this alert

On the right side of the upper pane, click the arrow to display more detailed information about the alert. Click it again to hide the detailed information.

To remove an alert, select it and click **Remove**.

To remove all alerts, click **Purge All**.

> **Note**: You can remove or purge only those alerts that have been resolved and informational alerts. Outstanding alerts remain in the list and on the **Dashboard** page until you resolve them.

Related Topics

- "ztC Endurance Console" on page 218

Support Logs Page

The **Support Logs** page enables you to generate diagnostic files, which include the ztC Endurance system's log files and configuration information at a particular moment in time. This information enables Penguin Solutions Customer Service to resolve an issue with the system.

For additional information, see:

- "Creating a Diagnostic File" on page 247

- "Deleting a Diagnostic File" on page 248

Related Topics

- "ztC Endurance Console" on page 218

Creating a Diagnostic File

Diagnostic files provide the ztC Endurance system's log files and configuration information at a particular moment in time. You create a diagnostic file to help Penguin Solutions Customer Service resolve issues with the system. You can also download it.

> **Note**: Automated Uptime Layer with Smart Exchange software allocates a fixed amount of storage space for diagnostic files. If a diagnostic file cannot be generated due to insufficient space, the system deletes previously created files.

To create diagnostic files

1. In the left-hand navigation panel, click **Support Logs** under **Alerts & Logs**.

2. On the **Support Logs** page, click **Generate Diagnostic File**.

After the diagnostic file has finished generating (as indicated by **complete** appearing in the **Status** column), you can download the file.

To download a diagnostic file

1. In the left-hand navigation panel, click **Support Logs** under **Alerts & Logs**.

2. On the **Support Logs** page, click anywhere in the row of the diagnostic file that you want to download. When a drop-down menu appears, click **Download** (⬇).

Related Topics

"ztC Endurance Console" on page 218

Deleting a Diagnostic File

Delete a diagnostic file from the ztC Endurance system.

To delete a diagnostic file

1. In the left-hand navigation panel, click **Support Logs** under **Alerts & Logs**.

2. On the **Support Logs** page, click anywhere in the row of the diagnostic file that you want to delete. When a drop-down menu appears, click **Delete**.

> **Note**: The drop-down menu does not appear while a diagnostic file is still generating.

Related Topics

- "ztC Endurance Console" on page 218

Configuring e-Alerts

Configuring email notifications (e-Alerts) enables the ztC Endurance system to send email to the configured list of recipients whenever an alert is generated or resolved. An e-Alert is sent every time an alert changes its state.

Prerequisite: In order for e-Alerts to function properly, you must configure the mail server. See "Configuring the Mail Server" on page 227.

To enable e-Alerts

1. In the left-hand navigation panel, click **e-Alerts** under **Notification**.

2. On the **e-Alerts** page, click the **Enable e-Alerts** box. The **List of Recipients** setting appears. Enable e-Alert delivery by specifying a valid sender's email address in the **List of Recipients** field. Specify email addresses for additional recipients on separate lines.

   Any email address that the SMTP server accepts is sufficient.

3. Click **Save**.

Notice: If you notice that delivery of e-Alerts seems delayed and you see the message `Message submission rate for this client has exceeded the configured limit` in the `sra-mgmt` log, the number of e-Alerts sent in a short period of time may have exceeded the email limits set in your mail server. Contact your email administrator if you need to increase the mail server's email limits.

Note: When you enable or update the e-Alert configuration, generate a test alert to confirm that the email messages are received.

To generate a test alert

Click **Generate Test Alert**. The Automated Uptime Layer with Smart Exchange software generates a test alert and sends a sample email with the subject "Test Alert" to all email recipients; SNMP sends traps to recipients of SNMP traps, if configured (see "Configuring SNMP Settings" on page 250); and Support Configuration sends a notification to your authorized Penguin Solutions service representative, if configured (see "Configuring Remote Support Settings" on page 264). A few minutes later, the software clears the test alert and sends an email notifying all email recipients that the alert has been cleared. View the **Alert History** page (see "Alert History Page" on page 246) for delivery status.

Notice: If the list of recipients changes immediately after a test alert is being generated, the recipient of the test alert may not receive the follow-up email stating that the alert has been cleared. To avoid this situation, do not change the list of recipients until after the recipient receives the follow-up email.

Related Topics

- "ztC Endurance Console" on page 218

Managing SNMP

Simple Network Management Protocol (SNMP) is a standard protocol for receiving alarms, sending traps, and monitoring system status. SNMP draws upon system-defining information that is stored in hierarchically configured management information bases (MIBs).

To configure a system to use SNMP, see "Configuring SNMP Settings" on page 250.

You can download a copy of the MIB file from the **Penguin Solutions Customer Service Portal** at https://service.stratus.com. For instructions for downloading software, see "Getting Help" on page 295.

Related Topics

- "ztC Endurance Console" on page 218

Configuring SNMP Settings

Configure Simple Network Management Protocol (SNMP) settings for the ztC Endurance system to allow SNMP management applications to remotely monitor your systems. (SNMP information pertains to systems as well as many system components, such as compute modules, storage modules, disks, and so on.) You can enable SNMP requests and SNMP traps:

- SNMP request—A request sent to the system to retrieve the values of objects listed in the Management Information Bases (MIBs) supported by the Automated Uptime Layer with Smart Exchange software. These MIBs include a system-specific MIB that is a collection of objects describing the ztC Endurance system. You can download a copy of the MIB file from the **Penguin Solutions Customer Service Portal** at https://service.stratus.com. For instructions for downloading software, see "Getting Help" on page 295.

- SNMP trap—A message initiated by one of the components in the ztC Endurance system after an event such as an alert that is then sent to an identified list of recipients (typically, a network management station (NMS)).

Managing Firewall Settings

You may need to modify some firewalls in order to communicate SNMP requests or traps from or to an NMS.

- When you installed the ztC Endurance software, your organization's firewall was automatically configured to open port 161 (Get/Walk) and 162 (Traps) for inbound and outbound traffic to the host . No user action is needed.

- If there is a firewall between the network that the ztC Endurance system is on and the location of the NMS, you may need to update your organization's firewall. Contact your network administrator about modifying the firewall.

Follow the appropriate procedure to enable SNMP requests or traps.

To enable SNMP requests

To enable SNMP requests, perform one of the following actions:

- Enable SNMP requests from the **SNMP Configuration** page of the ztC Endurance console:

    - Add an SNMPv3 user who can enable SNMPv3 requests and who has read-only access to the full MIB in the ztC Endurance system.

    - Configure access control for SNMPv1 and SNMPv2 requests, where you allow no users (**Restricted**) or any user using the default public community (**Unrestricted**) to send requests.

- Customize SNMP request functionality by editing `snmpd.conf` files. You can customize access control for SNMPv1 requests and SNMPv2 requests. You can also customize the list of users for SNMPv3 requests. For information, see To customize SNMP request functionality (below).

To enable SNMP requests from the SNMP Configuration page

1. In the left-hand navigation panel, click **SNMP Configuration** under **Notification**.

2. On the **SNMP Configuration** page, activate the check box next to **Enable SNMP Requests**.

3. The **List of Users for SNMP Requests (version 3)** and a read-only display of the `snmpd.conf` file appear.

    If a username appears below the **List of Users for SNMP Requests (version 3)**, the user's security level is displayed and a read-only display of the `snmpd.conf` file also appears. The user has read-only access to the full MIB.

    If a username does not appear, you can add an SNMPv3 user.

> **Notice**: Always use this procedure to add or remove SNMPv3 users. Do not edit the `/etc/snmp/snmpd.conf` and `/var/lib/net-snmp/snmpd.conf` files for this purpose.

To add an SNMPv3 user

a. Click the ➕ **Add** button, which opens the **Add a User** wizard.

b. Enter values for the following:

Username—The name of a user who has access to the SNMPv3 agent. The name must be unique.

Security Level—The user's security level. Valid values are:

- **No Authentication and No Privacy**: No security is applied to messages; messages are not authenticated or encrypted.

- **Authentication and No Privacy**: Messages are authenticated but not encrypted. You must enter values for **Authentication Type** and **Authentication Password**.

- **Authentication and Privacy**: Messages are authenticated and encrypted. You must enter values for **Authentication Type**, **Authentication Password**, **Encryption Type**, and **Encryption Password**.

When the security level includes authentication or privacy, the following fields appear:

Authentication Type—The user's type of authentication. Valid values are:

- **MD5**: Configure the message digest algorithm (MD5) as the user's authentication type.

- **SHA**: Configure the secure hash algorithm (SHA) as the user's authentication type.

Authentication Password—The user's required password, which is used to generate the secret authentication key. The password must be a minimum of eight characters and cannot be the same as the username.

Encryption Type—The user's type of encryption. Valid values are:

- **AES**: Configure the advanced encryption standard (AES) as the user's encryption type.

          ◦ **DES**: Configure the data encryption standard (DES) as the user's encryption type.

          Encryption Password—The user's required password, which is used to generate the secret encryption key. The password must be a minimum of eight characters and cannot be the same as the username.

    c. Click **Save** to save the changes.

4. Select an access option:

   Restricted (the default)—Prevents users from sending SNMPv1 requests and SNMPv2 requests.

   Unrestricted—Allows any user using the default public community to send SNMPv1 requests and SNMPv2 requests.

   Customized (available when `snmpd.conf` has been manually edited by a user; see [To customize SNMP request functionality](#), below)—Allows customized access.

5. Click **Save**. (Or click **Cancel** to restore the previously saved values.)

To customize SNMP request functionality by editing `snmpd.conf` files

Customize SNMP request functionality by editing `snmpd.conf` files.

Customize access control for SNMPv1 requests and SNMPv2 requests by editing the `/etc/snmp/snmpd.conf` file:

1. Log in to the host.

2. Manually edit the standard `/etc/snmp/snmpd.conf` file.

3. Save the file.

> **Notice**: The system-management software automatically restarts **snmpd** when it detects that **snmpd.conf** has changed; do not manually restart **snmpd**.

To enable SNMP traps

> **Notes**:
>
> 1. When you add a recipient for **SNMP Traps (version 3)**, you need to confirm that the engine ID of the trap user on the recipient server is 0x80001370017F000001.
>
> 2. When you enable or modify the SNMP trap settings, generate a test alert to confirm that traps are received.

1. In the left-hand navigation panel, click **SNMP Configuration** under **Notification**.

2. On the **SNMP Configuration** page, activate the check box next to **Enable SNMP Traps**.

3. Type the name of the SNMP **Community**, or keep the default (**public**).

4. Below the **List of Recipients of SNMP Traps (version 3)** is a list of the trap users, and the IP address of the recipient server where the trap user exists. The ztC Endurance system sends SNMPv3 traps to the trap user on the recipient server. Add a recipient, if necessary. Note that the system supports only one trap-recipient user.

To add a recipient

a. Click the **Add** button, which opens the **Add a Recipient** wizard.

b. Enter values for the following:

Recipient Address—The host name or the IPv4 address of the recipient server.

Username—The name of a trap user on the recipient server. The name must be unique for the recipient.

Security Level—The user's security level. Valid values are:

- **No Authentication and No Privacy**: No security is applied to messages; messages are not authenticated or encrypted.

- **Authentication and No Privacy**: Messages are authenticated, but not encrypted. You must enter values for **Authentication Type** and **Authentication Password**.

- **Authentication and Privacy**: Messages are authenticated and encrypted. You must enter values for **Authentication Type**, **Authentication Password**, **Encryption Type**, and **Encryption Password**.

When the security level includes authentication or privacy, the following fields appear:

Authentication Type—The user's type of authentication. Valid values are:

- **MD5**: Configure the message digest algorithm (MD5) as the user's authentication type.

- **SHA**: Configure the secure hash algorithm (SHA) as the user's authentication type.

Authentication Password—The user's required password, which is used to generate the secret authentication key. The password must be a minimum of eight characters and cannot be the same as the username.

Encryption Type—The user's type of encryption. Valid values are:

        ◦ **AES**: Configure the advanced encryption standard (AES) as the user's encryption type.

        ◦ **DES**: Configure the data encryption standard (DES) as the user's encryption type.

    Encryption Password—The user's required password, which is used to generate the secret encryption key. The password must be a minimum of eight characters and cannot be the same as the username.

    c. Click **Save** to save the changes.

5. Click **Save**. (Or click **Cancel** to restore the previously saved values.)

6. Click the **Generate Test Alert** button at the bottom of the configuration window.

The Automated Uptime Layer with Smart Exchange software generates a test alert and SNMP sends traps to recipients of SNMP traps; e-Alerts send a sample email with the subject "Test Alert" to all email recipients of e-Alerts, if configured (see "Configuring e-Alerts" on page 248); and Support Configuration sends a notification to Penguin Solutions Customer Service, if configured (see "Configuring Remote Support Settings" on page 264). Watch the Alert History log (see "Alert History Page" on page 246) for delivery status.

Related Topics

- "ztC Endurance Console" on page 218

- "Managing SNMP" on page 250

Configuring OPC Settings

Configure Open Platform Communication (OPC) settings to enable OPC server functionality, which publishes ztC Endurance system performance data for an OPC client to receive and display. This allows you to monitor the ztC Endurance system alongside other industrial equipment.

In order to use OPC functionality, you must install OPC client software (of your choice) on a separate computer and then configure the OPC client (see *"To install and configure an OPC client" on page 259*). The OPC client must be configured to receive data from the port on the ztC Endurance system that you configure for OPC. The default port is 4840, but you can specify another port number.

To configure OPC settings

1. In the left-hand navigation panel, click **OPC Configuration** under **Notification**.

2. On the **OPC Configuration** page, activate the check box next to **Enable OPC Server**.

3. The default port number is **4840**. If necessary, specify another number.

4. Activate either or both of the following check boxes, as appropriate for your system:

   - **Allow anonymous OPC client connections**: OPC clients do not need to provide a username and password when connecting to the OPC server on the ztC Endurance system. (When the check box is not activated, OPC clients must provide a username and password.)

   - **Allow OPC client connections using usernames and passwords configured from Users**: OPC clients can connect to the OPC server on the ztC Endurance system with the same username and password used to log in to the ztC Endurance console. (When the check box is not activated, OPC clients cannot log in to the OPC server using local-user account usernames and passwords, as specified on the **Users** page. See "Managing User Accounts" on page 245.)

5. Click **Save**. (Or click **Cancel** to restore the previously-saved values.)

To install and configure an OPC client

You must install OPC client software and configure an OPC client on a separate computer. You can choose OPC client software from the many versions that exist in the marketplace. The following procedure describes how to install and configure an OPC client using UaExpert® software from Unified Automation.

Installing and Configuring an OPC Client Using UaExpert Software

> **Note**: In addition to reading the following procedure, follow instructions with the UaExpert software.

1. Download and install the Linux version of the UaExpert software. See https://www.uni-fied-automation.com/products/development-tools/uaexpert.html.

2. If starting UaExpert software for the first time, follow the instructions with the software for an initial start-up.

3. Run the UaExpert software.

   The **Unified Automation UaExpert - The OPC Unified Architecture Client - NewProject** main window opens.

4. In the menu bar, click **Server** and select **Add**. The **Add Server** dialog box appears.

5. Click the **Advanced** tab.

6. In the **Endpoint Url** box, enter the URL of the endpoint, which is the ztC Endurance system's cluster IP address (for example, **opc.tcp://***tcp_cluster_ip_address***:4840/**).

7. For **Security Settings**, select **None** for both **Security Policy** and **Message Security Mode**.

8. For **Authentication Settings**, select one of the following, as required for your configuration:

   - **Anonymous**: Select if you activated the check box for allowing anonymous OPC client connections.

   - **Username** and **Password**: Enter values if you activated the check box for allowing OPC client connections using usernames and passwords. The username and password you enter must be identical to the username and password that you add for a read-only user on the ztC Endurance system for

OPC access. See "Managing User Accounts" on page 245 for information on adding a user to the ztC Endurance system.

9. Click **OK** to close the **Add Server** dialog box.

The main window reappears. In the left panel, the name of the server appears in the **Project** box, under **Servers**.

10. Select the new server and then click the connect button, which appears in the tools bar to the right of the minus-sign icon.

When the client connects successfully to the server, the **Address Spaces** box of the main window displays the end point of the server.

In the **Address Spaces** box, you can click the top level to expand and explore the available data values. In the **Attributes** box, the **Value** column displays the current value of the selected item.

Related Topics

- "Displaying OPC Output" on page 260

- "ztC Endurance Console" on page 218

- "Using the ztC Endurance Console" on page 217

Displaying OPC Output

After you have enabled OPC server functionality on the ztC Endurance system, an OPC client, located on a computer that is not the ztC Endurance system, can publish the system's performance data. The data is divided into address spaces, where each address space is divided into sub-objects that contain several data items. The OPC server functionality on the ztC Endurance system passes values for the data items to the OPC client, which then displays the data.

This topic describes how to display a ztC Endurance system's basic inventory and health information using an OPC client created with UaExpert® software from Unified Automation.

For more information about navigating the OPC client, click the **Help** tab in the navigation menu at the top of the client window.

To display OPC output

1. On the computer where you have created an OPC client, open the client.

2. In the (left) **Project** panel, click **Project**, **Servers**, and then **ztC OPC Server**.

After you select **ztC OPC Server**, the **Address Space** panel under **Project** displays the **Root** hierarchy.

3. In the **Address Space** panel, under **Root**, click **Objects**. Under **Objects**, you can select **Server** and **ztCEndurance**. Each physical component of the system is represented by a data item in the panel; click the data item to check the component's status. The **Attributes** panel on the right side of the window displays the value of the selected data item.

The following tables describe some of the data items that can be viewed on the OPC client.

Server

For information about the compute module that is currently running the OPC server, click **Server** and then **ServerStatus** to view the **BuildInfo** sub-object.

The **BuildInfo** sub-object displays values for the following data items:

Data Items	Description
ProductUrl	Displays **http://www.stratus.com/**.
ManufacturerName	Displays **Stratus Technologies Ireland, Ltd**.
ProductName	Displays the product name of the hardware (**ztC Endurance**).
SoftwareVersion	Displays the version number of the Automated Uptime Layer with Smart Exchange software.
BuildNumber	Displays the Automated Uptime Layer with Smart Exchange software build number.
BuildDate	Displays the date of the Automated Uptime Layer with Smart Exchange software build.

For additional information about the **Server** object, see *Part 5: Information Model* of the *OPC Unified Architecture Specification*, which is available at opcfoundation.org.

ztCEndurance

The **ztCEndurance** object divides the address space according to each of the physical components of the connected ztC Endurance system. The sub-objects are as follows: **Compute Modules**, **I/O Modules**, **Mirrors**, **Power Supply Units**, **Storage Modules** and **System**. These sub-objects can be opened to display their components.

For example, the data item **ComputeModuleA** expands to display the following data items:

Data Items	Description
BMCIPv4AddressA	Displays the IP address of BMC A. Data type: string
BMCIPv4AddressB	Displays the IP address of BMC B. Data type: string
Healthy	Indicates whether the module is healthy: *true* indicates the module is present, running green-checked, and is not in maintenance mode; *false* indicates that the module is not present, is not running green-checked, and/or is in maintenance mode. Data type: Boolean
Present	Indicates whether the module is present; *true* or *false* value. Data type: Boolean
State	Indicates whether the module is online. Example value: *ONLINE/ACTIVE*. Data type: string

Additional Data Items

The following table lists data items not found in the previous example, but that appear in other sub-objects under the **ztCEndurance** tab.

Data Items	Description
LinkSpeed	Numeric value that indicates the link speed of the management PC. Data type: Int64
LinkUp	Indicates that the network wire is connected and communicating; *true* or *false* value. Data type: Boolean
EnduranceRemaining	Percentage value that indicates how much of the system's predicted life span remains. Data type: UInt64

System

The data items for the system provide high-level status information as well as information about access methods for the overall system.

Data Items	Description
Healthy	Indicates the system health with *true* or *false* value. Data type: Boolean
Model	Displays the system model name. Data type: string
OutstandingSeverityLevel	Corresponds to the overall system status icon on the login page. Data type: UInt32
StratusManagementURL	Displays the HTTPS URL of the ztC Endurance system. Data type: string
SecureManagementURL	Displays the HTTPS URL for the ztC Endurance console. Data type: string

Related Topics

- "Configuring OPC Settings" on page 257

Configuring Remote Support Settings

When you log on to the ztC Endurance system for the first time, configure support configuration settings that enable the ztC Endurance system to send support notifications (alerts) to Penguin Solutions Customer Service via the Active Service Network (ASN) when an event requires attention.

To configure support configuration settings

> **Note**: When you enable or modify settings for **Enable Support Notifications**, generate a test alert to confirm that Penguin Solutions Customer Service can receive system health messages from your system.

1. In the left-hand navigation panel, click **Support Configuration** under **Remote Support**.

2. On the **Support Configuration** page, modify the settings, as appropriate for your system:

> **Notice**: If you plan to enable remote support access, the ASNConnect by BeyondTrust software must already be installed. The **Enable Remote Support Access** option will be grayed out if it is not installed.

- **Enable Remote Support Access** allows Penguin Solutions Customer Service to remotely connect to the ztC Endurance system for troubleshooting purposes. You can enable and then disable this setting, as needed.

- **Enable Support Notifications** allows the ztC Endurance system to send health and status notifications to Penguin Solutions Customer Service. **Enable Support Notifications** sends an alert for any event that requires attention. You can also select **Enable Periodic Reporting** to send a periodic "heartbeat" call-home message to Penguin Solutions Customer Service.

> **Note**: To allow Penguin Solutions Customer Service to monitor your system for failed ASN heartbeat alarms, you must enable heartbeats in the **Penguin Solutions Customer Service Portal**. See "To configure ASN "heartbeat" alarms" on page 267.

3. Click **Save** (or click **Cancel** to restore the previously saved values).

4. Configure your organization's firewall to allow support messages.

To configure your firewall to allow support messages

Use the following information as needed to configure your organization's firewall to allow communication with Penguin Solutions Customer Service:

Message Type: Call-Home and Licensing
Protocol: TCP
Port: 443
Stratus support server address: *.stratus.com

Message Type: Support Diagnostics
Protocol: TCP
Port: 443
Stratus support server address: *.stratus.com

Message Type: ASNConnect
Protocol: TCP
Port: 443, Default proxy port: 3128 (You can change the default proxy port number.)
Stratus support server address: *.asnconnect.stratus.com

To enable SNMP management systems to receive alerts and send traps to the ztC Endurance system, configure the firewall for the following:

Message Type: SNMP
Protocol: SNMP
Port: 161 (Get/Walk) 162 (Traps)

5. Generate a test alert.

To generate a test alert

Click **Generate Test Alert**. The system generates a test alert and Support Configuration sends a notification to Penguin Solutions Customer Service; e-Alerts send a sample email with the subject "Test Alert" to all email recipients of e-Alerts, if configured (see "Configuring e-Alerts" on page 248); and SNMP sends traps to recipients of SNMP traps, if configured (see "Configuring SNMP Settings" on page 250). View the **Alert History** page (see "Alert History Page" on page 246) for delivery status. A subsequent alert is generated if the support notification fails.

To configure ASN "heartbeat" alarms

1. On the **Support Configuration** page, ensure that **Enable Support Notifications** is selected, and select **Enable Periodic Reporting**. Click **Save**.

2. Log in to the **Penguin Solutions Customer Service Portal** and open the ztC Endurance asset that you want to enable for periodic ASN "heartbeat" alarms. Click the **Service Network Configuration** link.

3. On the **Service Network Configuration** page, select **Enable Heartbeat**. For **Heartbeat Interval**, select **Days** and **1**, then select **Default Heartbeat Grace Period**. Click **Save Configuration**.

Related Topics

- "ztC Endurance Console" on page 218

- "Getting Help" on page 295

Configuring Internet Proxy Settings

Configure proxy settings for the ztC Endurance system if your organization requires a proxy server to access the Internet and you have a service agreement with Penguin Solutions or another authorized ztC Endurance service representative.

A proxy server provides a secure bridge between the ztC Endurance system and the Internet. Automated Uptime Layer with Smart Exchange software uses proxy server information for only outbound HTTP traffic related to support notification messaging and remote support access features.

To configure Internet proxy settings

1. In the left-hand navigation panel, click **Proxy Configuration** under **Remote Support**.

2. On the **Proxy Configuration** page, click the **Enable Proxy** box to enable proxy service.

3. For **Scheme**, select either **https** or **http**.

4. In the **Proxy Server** box, type the fully-qualified proxy server host name or IP address.

5. In the **Port Number** box, type the port number if it is different from the default number (3128).

6. If the proxy server requires authentication, click the **Enable Authentication** box and type the **Username** and **Password**.

   If you do not type a password, the previous password continues to be required. If the previous password was empty and you do not enter a new password, the password remains empty.

7. Click **Save** (or click **Cancel** to restore the previously saved values).

Related Topics

- "ztC Endurance Console" on page 218

Upgrade Kits Page

The ztC Endurance **Upgrade Kits** page enables you to upload and manage upgrade kits. Use this page to upgrade the system to newer versions of the Automated Uptime Layer with Smart Exchange software.

For complete instructions on downloading upgrade kits from the **Penguin Solutions Customer Service Portal** and installing them on the system, see "Upgrading the System Software" on page 188.

The **Upgrade Kits** page supports the following functions:

- **Qualify** – When an upgrade kit in the **Ready** state is selected, the **Qualify** button verifies that no version conflicts exist, and that the system is in a healthy state for the upgrade to start. If qualification fails, resolve the issue described in the error message.

- **Upgrade** – When an upgrade kit in the **Ready** state is selected, the **Upgrade** button starts the upgrade process.

- **Delete** – The **Delete** button deletes the selected upgrade kit from the system.

- **Add** – The **Add** button loads an upgrade kit into the **Upgrade Kits** page so that you can apply the upgrade. Before adding an upgrade kit file, you must first download it from the **Penguin Solutions Customer Service Portal** to a local system.

Upgrade Kit States

The left column of the **Upgrade Kits** page displays the state of each kit file that has been uploaded to the system. Possible kit states are as follows:

- **Ready** – Kit is ready to apply.

- **Applied** – Kit has been successfully applied.

- **Apply in progress** – Kit is being applied. This state also displays the status of the host, compute module A and compute module B.

- **Validating** – Kit checksum is being verified to determine if the kit file was corrupted.

- **Unsupported** – Kit is older than the version currently running on the system. The **Upgrade** button is unavailable for kits that are unsupported, because the system does not support installing older versions.

- **Broken** – Checksum verification failed.

Related Topics

"Upgrading the System Software" on page 188

"ztC Endurance Console" on page 218

"Using the ztC Endurance Console" on page 217

8

Chapter 8: Overview of the BMC Web Console

The BMC Web console is a Web-based interface that allows system administrators and authorized users to manage aspects of the ztC Endurance system. The BMC Web console supports the following capabilities for the user:

- Remote power control

- Remote keyboard/video/mouse (KVM)

- Remote media

- Passive examination of logs and sensors for diagnostic purposes

> **Note**: Each BMC only provides information available to its own compute module. You must query both BMCs to obtain all information.

Related Topics

- "Prerequisites for Using the BMC Web Console" on page 271

- "Connecting to the BMC Web Console" on page 272

- "Logging In to the BMC Web Console" on page 273

- "Using the BMC Web Console Online Help" on page 274

- "Unsupported BMC Web Console Features" on page 275

Prerequisites for Using the BMC Web Console

Before connecting to the BMC Web console, make sure your management PC adheres to the following requirements:

Required Browser Settings

- Allow pop-ups from this site

- Allow file downloads from this site

- Enable JavaScript for this site (any JavaScript release can be used)

- Enable cookies for this site

The BMC Web console supports the following browsers:

- Firefox® Version 2.0 and above (Linux® and Windows®)

- Chrome™ Version 2.0 and above (Windows)

- Microsoft® Edge Version 105 and above (Windows)

Additional Requirements

For remote media connections, reduce the link speed to 100 Mbps in the network interface of the system that is hosting the remote media.

Related Topics

- "Overview of the BMC Web Console" on page 271

Connecting to the BMC Web Console

You connect to the BMC Web console by entering the IP address of the BMC Web console into a supported Web browser. For more information on supported browsers, see "Prerequisites for Using the BMC Web Console" on page 271. You can obtain the IP address from the ztC Endurance console. If the ztC Endurance console is not running, you can obtain the IP address from the BIOS setup utility.

To obtain the BMC Web console IP address from the ztC Endurance console

1. Log in to the ztC Endurance console.

2. In the left-hand navigation panel, click **Hardware** under **System**.

3. Under **Hardware**, select the correct compute module from the list.

4. Click the **BMC** tab in the bottom pane to display the IP addresses.

5. Click the IP address of the BMC Web console to which you want to connect.

To obtain the BMC Web console IP address from the BIOS setup utility

1. Enter the BIOS setup utility.

2. Click **Server Mgmt**.

3. Under **Server Mgmt**, click **BMC Network Configuration**.

4. Scroll to **LAN Channel 1** or **LAN Channel 8** and locate **Station IP Address**. Copy the IP address, paste it into your Web browser, and hit **Enter** to connect to the BMC Web console.

Next Steps

After you connect to the BMC Web console IP address, log in as described in "Logging In to the BMC Web Console" on page 273.

Related Topics

- "ztC Endurance Console" on page 218

- "Overview of the BMC Web Console" on page 271

Logging In to the BMC Web Console

After you connect to the BMC Web console, log in to it by entering an authorized username and password on the login page. When logging in to the BMC Web console for the first time, use the default credentials:

Username: **admin**

Password: **admin**

> **Notices**:
>
> 1. Using the BMC Web console to change the **admin** password after system installation results in system failure.
>
> 2. Prior to system installation, the BMC Web console username and password are set to the default **admin**. During the installation, you are asked to provide a new password for the **admin** account. After the installation is complete, you must create a separate user account for personal use. Passwords must be 8-16 characters. For more information, see "To create a new user account " on page 274

To log in to the BMC Web console

1. Next to **Username** on the BMC login page, type your username (**admin**, if this is your first login).

2. Next to **Password**, type your password (**admin**, if this is your first login).

3. Click **Sign me in**.

> **Note**:
> While multiple users can be logged in to the BMC Web console, only one user can connect to the H5Viewer (the KVM) at a time. If a new user attempts to log in, the current user receives a message in the H5Viewer asking permission to give the new user access. If the current user does not respond, the new user is automatically granted access.

To create a new user account

1. From the BMC Web Console dashboard, click **Settings**.

2. Under **Settings**, click **User Management**.

3. Select the desired channel from the drop-down button at the top of the page. Click any of the available slots displayed.

4. Enter the new user's information and click **Save** when finished.

Next Steps

If you are performing a system software installation, return to the next step of "To complete the installation" on page 161.

Related Topics

- "Overview of the BMC Web Console" on page 271

Using the BMC Web Console Online Help

The BMC Web console contains a built-in help system that provides information about each page.

The Help button ( ) is located in the top right corner of each page. When you click the Help button, information about the page appears. Click the Help button again to hide the information.

Related Topics

- "Overview of the BMC Web Console" on page 271

Configuring DNS Settings in the BMC Web Console

Configure Domain Name System (DNS) settings in the BMC Web console to register your system's name with your DNS server while in DHCP mode. This mode is called Dynamic DNS (DDNS).

The BMC Web console has three DDNS configuration options:

- **nsupdate** (the default value) communicates directly with the DNS nameserver.

- **FQDN** requests that the DHCP server set the FQDN in the DNS server as specified by the client.

- **Hostname** requests that the DHCP server set the hostname in the DNS server.

To specify the DDNS method in the BMC Web Console

1. [Log in to the BMC Web console](#).

2. Under **Settings** in the left-hand navigation panel, click **Network Settings**, then click **DNS Configuration** to open the **DNS Configuration** page.

3. Select the desired configuration, or contact your network administrator for the recommended settings.

4. Click **Save**.

Related Topics

- "Overview of the BMC Web Console" on page 271

Unsupported BMC Web Console Features

Some actions within the BMC Web console are not supported for customer use.

> **Notice**: Performing any function that does not appear in "Overview of the BMC Web Console" on page 271 puts the operation of the ztC Endurance system at risk.

Browser Functions to Avoid

Once you are logged in to the BMC Web console, avoid using the following browser functions:

- The back and forward buttons on the browser

- The refresh button on the browser

- The F5 key on the keyboard

- The backspace key on the keyboard

- The number pad on the keyboard

Instead of using the preceding browser functions, use the buttons located in the user menu in the top right corner of the BMC Web console.

Avoid Using Caps Lock Key in KVM

The BMC KVM interface does not correctly interpret the Caps Lock key. Always use the Shift key when typing capital letters.

Do Not Use BMC Web Console to Power Off Running System

Use the ztC Endurance console, **not** the BMC Web console, to power off a running ztC Endurance system.

Using the BMC Web console to power off a system may result in the following undesirable outcomes:

- If a user inadvertently powers off the active compute module, application downtime can occur.

- If a user powers off the standby compute module, the system-management software attempts to power it back on in order to maintain maximum system availability.

Part 3: ztC Endurance Supporting Documents

See the following support documents for release information, and for reference and troubleshooting information.

-
-

Related Topics

-
-

9

Chapter 9: Release Notes: Automated Uptime Layer with Smart Exchange Release 2.0.0.0 (Linux)

These Release Notes (updated 3:13 PM EST on 7/8/2025) are for Automated Uptime Layer with Smart Exchange Release 2.0.0.0. See the following sections:

New Features and Enhancements

Stratus Technologies is now part of Penguin Solutions. To learn more about Penguin Solutions, visit https://www.penguinsolutions.com.

Automated Uptime Layer with Smart Exchange Release 2.0.0.0 introduces the following ztC Endurance systems running on Red Hat Enterprise Linux (RHEL) version 9.4 and Automated Uptime Layer with Smart Exchange:

- ztC Endurance 3100 and 3110

- ztC Endurance 5100 and 5110

- ztC Endurance 7100 and 7110

- ztC Endurance 9110

See for more information.

Important Considerations

KVM Not Supported in Automated Uptime Layer with Smart Exchange Release 2.0.0.0

Kernel-based Virtual Machine (KVM) functionality is not supported in Release 2.0.0.0.

Do Not Update the Ubuntu System Software

Do not enable Ubuntu unattended upgrades or update the Ubuntu system software of the Standby OS from any source other than Stratus unless instructed to do so. To prevent incompatibilities and ensure the proper fault-tolerant operation of your system, use only the release and software packages that are installed with the Automated Uptime Layer with Smart Exchange software. All supported Ubuntu package updates will be bundled in future releases.

Do Not Change BIOS Settings

Do not change any BIOS settings from their default settings without first being told to do so by Penguin Solutions Customer Service.

Known Issues

Missing Management IP Addresses After `sra-sysconfig.exe` Runs

If network changes during system configuration take longer than expected to take effect, the `sra_sysconfig.exe` script may not display any IP addresses after printing the message `Management Addresses for ztC Endurance console.`

The workaround is to wait two minutes, then run the following command to ensure the addresses are available:

```
/opt/stratus/sbin/zencli mgmt_addrs
```

Run Installation Locally When Setting Management Network IPs to Static Addresses

When performing a ztC Endurance installation that changes the management network IPs to static addresses, perform the installation from the BMC Web console Remote Control H5Viewer (the KVM), or from the monitor and keyboard connected directly to the ztC Endurance system. Otherwise, the installation may not complete successfully.

Possible Message About Duplicate IPv4 Address When `snmpd.service` Starts

When the Simple Network Management Protocol (SNMP) Daemon is started, the following error message may appear in the messages log:

```
DuplicateIPv4 address detected, some interfaces may not be visible in IP-
MIB
```

This error message is harmless and can be safely ignored. For more information, see the RHEL documentation at https://access.redhat.com/solutions/3660111.

Embedded Network Interface Requires Manual Activation After Smart Exchange

After removing an I/O module, the embedded I210 network interface (eno1 or eno2) is disconnected. If a Smart Exchange occurs and the I/O module is then re-inserted, NetworkManager does not re-establish the connection for the embedded I210 network interface, and it no longer has an assigned IP address.

To fix this issue, manually activate the connection by running the following command from the host:

- If I/O module A is removed and re-inserted:

```
nmcli con up eno1
```

- If I/O module B is removed and re-inserted:

```
nmcli con up eno2
```

Restart NetworkManager After Adapter Replacement

In cases where a PCIe adapter is part of a bonded team, the adapter is not automatically added back to the bond after it is swapped out and replaced. The new PCIe adapter appears as **Online**, but ztC Endurance management reports its state as Simplex.

To re-add the new PCIe adapter to the bonded team, restart NetworkManager using the following command:

```
systemctl restart NetworkManager
```

Enable SNMP Requests to Fix Location of `snmpd.conf` File

On a newly installed system, the **Enable SNMP Requests** check box on the ztC Endurance console **SNMP Configuration** page is disabled by default. This results in the system's `/etc/snmp/snmpd.conf` file pointing to a non-existent file, preventing the RHEL SNMP daemon from working correctly.

To fix this problem, activate the **Enable SNMP Requests** check box on the ztC Endurance console **SNMP Configuration** page.

802.3ad Bond Mode Requires Setting lacp_rate to Fast to Avoid Delays

When using 802.3ad bonding mode, ping delays of 10-12 seconds can occur.

To increase response time, change `lacp_rate` from **slow** (the default) to **fast** by manually editing the configuration file as follows:

1. Access the 802.3ad bonding team file: `/etc/NetworkManager/system-connections/bondN.nmconnection`.

2. Add the line `lacp_rate=fast` under the `[bond]` section:

   ```
   [bond]

   lacp_rate=fast

   mode=802.3ad
   ```

3. Run the following commands, where N is the bond number, to have NetworkManager process the change:

   ```
   nmcli c reload
   nmcli c down bondN
   nmcli c up bondN
   ```

Static Link Aggregation Bonding Modes Encounter Delays

Adding a PCIe adapter to a team configured in static-bonding mode could result in long network delays when an I/O module is being brought up.

Certain bonding modes require that the physical switch be set up with a static link aggregation group. These bonding modes are: Balance-XOR, Balance-RR, BC, and Broadcast. However, I/O failovers may encounter delays of between 20 and 30 seconds in these conditions.

Avoid using these modes unless absolutely necessary. Alternate modes include: active-backup (default), balance-tlb, balance-alb, and 802.3ad.

Incorrect Suffix Added to Hostname After Successful Installation

After a successful ztC Endurance installation, hostnames may have an extra domain name appended, causing access failures. For example, the expected host system name is **XX.cdx.eng.stratus.com**, but it is incorrectly named **XX.cdx.eng.stratus.com.cdx.eng.stratus.com**.

To fix this issue, reset the hostname to the correct format using the following command:

```
hostnamectl set-hostname XX.cdx.eng.stratus.com
```

Verify the change by entering `hostname` or `hostnamectl`.

Do Not Modify Graphical Mode Install in Kickstart File

This note applies only if you plan to customize your RHEL installation by modifying the ztC Endurance kickstart file.

When making kickstart file modifications, do not make changes to the section `# === Use graphical mode install ===`. Graphic mode must remain set to `graphic`.

Do Not Use BIOS to Disable Wake On LAN

Setting `Wake On LAN` to `Disabled` in the BIOS does not disable this functionality. Instead, use the `ethtool` command with the `--wol` option to disable Wake On LAN. For more information, see the RHEL documentation at https://docs.redhat.com/en/documentation/red_hat_enterprise_linux/6/html/deployment_guide/s1-ethtool#s1-ethtool.

Install ASNConnect After You Replace a Compute Module

If you replace a compute module, you must contact your authorized Penguin Solutions service representative (see "Getting Help" on page 295) so they can remove the ASNConnect software associated with the old compute module. After that, you can install ASNConnect on the new compute module as described in "Enabling ASNConnect" on page 213.

IPv6 Must Be Enabled in BMC Web Console

ztC Endurance systems are shipped with the **Enable IPv6** setting selected in the BMC Web console. The ztC Endurance system does not work if you disable this setting.

The **Enable IPv6** setting is located in **Settings**>**Network Settings**>**Network IP Settings**. When you select this setting, the **Enable IPv6 DHCP** setting appears; select this setting as well. For more information about the BMC Web console, see "Overview of the BMC Web Console" on page 271.

Installation with IPv6 Static Addresses in the yaml File Failed to Set RHEL Host IPs

If you are using an installation configuration file (`zen_config.yaml`) that includes IPv6 static addresses, the ztC Endurance installation may fail with error messages similar to the following:

```
Unable to install ztC Endurance components on the RHEL host
```

Or:

```
Error: failed to modify ipv6.addresses: invalid prefix
```

These errors occur because the IPv6 prefix is incorrectly appended twice to the end of the address.

To avoid this issue, specify the IPv6 prefix only at the `IPv6 prefix` prompt, or in the yaml field that asks for the IPv6 prefix. Do not include the prefix at the `IPv6 address` prompt, or in the yaml field that asks for the IPv6 address.

If this issue does occur, manually set the necessary static IPV6 addresses on the host.

Reconfiguring System from Static IPv6 Addresses to Dynamic IPv4 Addresses

If your ztC Endurance system has been configured with static IPv6 addresses and you are reconfiguring the system to use dynamic (DHCP) IPv4 addresses, perform the following procedure before running `net_config.sh`:

1. Log in to one of the BMC Web consoles.

2. Select **Settings**, then select **Network Settings**, and then select **Network IP Settings**.

3. On the **Network IP Settings** page:

   a. Select the **Enable IPv6 DHCP** check box.

   b. Enter **0** in the **IPv6 Address** field.

   c. Click **Save** and exit the BMC Web console.

Repeat this procedure on the other BMC Web console.

Resolve Active Alerts Before Upgrading System Software

If you attempt to upgrade the system software and the ztC Endurance system has active alerts of any kind, including informational alerts, the upgrade fails with the following message:

    Upgrade cannot proceed until all active alerts have been resolved.

After you resolve all active alerts, you can proceed with the upgrade.

Use the `--forcerescan` Argument to View a Mapped LUN 0

Generally, you should not map a LUN 0. However, if you do attempt to map a LUN 0 and it is not visible in the output after you run `rescan-scsi-bus.sh`, you may need to add the argument `--forcerescan`.

The current LUN numbers appear as the last digit in the `host:channel:id:lun` SCSI array. In the following example, LUNs 1 and 2 already exist and are mapped, and a LUN 0 has been added to the array, but the LUN 0 does not appear when you run the `rescan-scsi-bus.sh`:

```
rescan-scsi-bus.sh
Scanning SCSI subsystem for new devices
Scanning host 0 for  all SCSI target IDs, all LUNs
```

```
sg0 changed:  device 0 0 0 0 ...
from:Enclosure 0 Channel: 00 Id: 00 Lun: 00
to: Direct-Access  TE  Model: 4006           Rev: I200
      Type:   Enclosure                      ANSI SCSI revision: 06


Scanning for device 0 0 0 1 ...
OLD: Host: scsi0 Channel: 00 Id: 00 Lun: 01
Vendor: SEAGATE  Model: 4006          Rev: I200
Type:   Direct-Access                 ANSI SCSI revision: 06
Scanning for device 0 0 0 2 ...
OLD: Host: scsi0 Channel: 00 Id: 00 Lun: 02
Vendor: SEAGATE  Model: 4006          Rev: I200
Type:   Direct-Access                 ANSI SCSI revision: 06
Scanning host 1 for  all SCSI target IDs, all LUNs
sg3 changed:  device 1 0 0 0 ...
from:Enclosure 1 Channel: 00 Id: 00 Lun: 00
to: Direct-Access  TE  Model: 4006           Rev: I200
      Type:   Enclosure                      ANSI SCSI revision: 06
Scanning for device 1 0 0 1 ...
OLD: Host: scsi1 Channel: 00 Id: 00 Lun: 01
      Vendor: SEAGATE  Model: 4006           Rev: I200
      Type:   Direct-Access                  ANSI SCSI revision: 06
Scanning for device 1 0 0 2 ...
OLD: Host: scsi1 Channel: 00 Id: 00 Lun: 02
      Vendor: SEAGATE  Model: 4006           Rev: I200
      Type:   Direct-Access                  ANSI SCSI revision: 06
0 new or changed device(s) found.
0 remapped or resized device(s) found.
0 device(s) removed.
multipath -ll
mpathb (3600c0ff0005d03e8c12e856701000000) dm-0 SEAGATE,4006
size=279G features='0' hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| `- 0:0:0:1 sda     8:0    active ready running
`-+- policy='service-time 0' prio=10 status=enabled
`- 1:0:0:1 sdc     8:32   active ready running
mpathc (3600c0ff0005d0637839fa76701000000) dm-1 SEAGATE,4006
size=186G features='0' hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
```

```
| `- 1:0:0:2 sdd     8:48    active ready running
`-+- policy='service-time 0' prio=10 status=enabled
`- 0:0:0:2 sdb     8:16    active ready running
```

To view the LUN 0 along with the other LUNs in the array, add the `--forcerescan` argument as shown in the following example:

```
rescan-scsi-bus.sh --forcerescan
Syncing file systems
Scanning SCSI subsystem for new devices and remove devices that have dis-
appeared
Scanning host 0 for  all SCSI target IDs, all LUNs
sg0 changed:  device 0 0 0 0 ...
from:Enclosure 0 Channel: 00 Id: 00 Lun: 00
to: Direct-Access  TE  Model: 4006            Rev: I200
REM: Host: scsi0 Channel: 00 Id: 00 Lun: 00    ANSI SCSI revision: 06
NEW: Host: scsi0 Channel: 00 Id: 00 Lun: 00
     Vendor: SEAGATE  Model: 4006              Rev: I200
     Type:   Direct-Access                     ANSI SCSI revision: 06
Scanning for device 0 0 0 1 ...
OLD: Host: scsi0 Channel: 00 Id: 00 Lun: 01
Vendor: SEAGATE  Model: 4006            Rev: I200
Type:   Direct-Access                   ANSI SCSI revision: 06
Scanning for device 0 0 0 2 ...
OLD: Host: scsi0 Channel: 00 Id: 00 Lun: 02
     Vendor: SEAGATE  Model: 4006              Rev: I200
     Type:   Direct-Access                     ANSI SCSI revision: 06
Scanning host 1 for  all SCSI target IDs, all LUNs
sg3 changed:  device 1 0 0 0 ...
from:Enclosure 1 Channel: 00 Id: 00 Lun: 00
to: Direct-Access  TE  Model: 4006            Rev: I200
REM: Host: scsi1 Channel: 00 Id: 00 Lun: 00    ANSI SCSI revision: 06
NEW: Host: scsi1 Channel: 00 Id: 00 Lun: 00
     Vendor: SEAGATE  Model: 4006              Rev: I200
     Type:   Direct-Access                     ANSI SCSI revision: 06
Scanning for device 1 0 0 1 ...
OLD: Host: scsi1 Channel: 00 Id: 00 Lun: 01
Vendor: SEAGATE  Model: 4006            Rev: I200
Type:   Direct-Access                   ANSI SCSI revision: 06
```

```
Scanning for device 1 0 0 2 ...
OLD: Host: scsi1 Channel: 00 Id: 00 Lun: 02
      Vendor: SEAGATE  Model: 4006           Rev: I200
      Type:   Direct-Access                  ANSI SCSI revision: 06
2 new or changed device(s) found.
[0:0:0:0]
[1:0:0:0]
0 remapped or resized device(s) found.
2 device(s) removed.
[0:0:0:0]
[1:0:0:0]
```

Do Not Modify Time Zone Except on RHEL Host

Other than on the RHEL host, do not modify the time zone on your ztC Endurance system. For example, do not attempt to modify the time zone on the Standby OS or in the BIOS.

Delay in Client/Server-Based Applications During Smart Exchange While Communicating Over IP Addresses on X710 PCIe Adapters

Any client/server-based application that communicates over an IP address that is assigned to a networked team configured over an X710 PCIe adapter may experience a 7- to 25-second delay during or after a Smart Exchange operation.

Accessing Stratus Knowledge Base Articles

For information about accessing Stratus Knowledge Base articles, see "Getting Help" on page 295.

Getting Help

If you have a technical question about ztC Endurance systems, see "Getting Help" on page 295.

10

Chapter 10: System Reference Information

See the following support documents for reference and troubleshooting information.

System Specifications

Compute Module (Quantities listed are for each compute module)	
Physical dimen-sions	Maximum weight: • 31x0: 10 lb (4.5 kg) • 51x0: 12 lb (5.4 kg) • 71x0: 12 lb (5.4 kg) • 9110: 12 lb (5.4 kg) Length: 22.25 in. (565.2 mm) Width: 8.625 in. (219 mm) Height: 1.625 in. (41.2 mm)
Processors	3100: One Intel® Sapphire Rapids 2.0 GHz 12-core processor

	3110: One Intel Emerald Rapids 2.4 GHz 12-core processor 5100: Two Intel Sapphire Rapids 2.0 GHz 12-core processors 5110: Two Intel Emerald Rapids 2.4 GHz 12-core processors 7100: Two Intel Sapphire Rapids 2.0 GHz 24-core processors 7110: Two Intel Emerald Rapids 2.2 GHz 28-core processors 9110: Two Intel Emerald Rapids 2.8 GHz 32-core processors
System memory	RDIMM capacity: 64 GB 31x0: 8 physical, DDR5 RDIMM slots. Available RDIMM configurations are 64 GB, 128 GB, or 256 GB. 51x0: 16 physical, DDR5 RDIMM slots. Available RDIMM configurations are 128 GB, 256 GB, or 512 GB. 71x0: 16 physical, DDR5 RDIMM slots. Available RDIMM configurations are 256 GB, 512 GB, or 1024 GB. 9110: 16 physical, DDR5 RDIMM slots. Available RDIMM configuration is 1024 GB.
Ports	One UART (USB-C) debug port. For use by Penguin Solutions Customer Service. One VGA port Two USB 3.0 ports
Storage Module	
Physical dimensions	Maximum weight: 10 lb (4.5 kg) Length: 22.25 in. (565.2 mm) Width: 8.625 in. (219 mm) Height: 1.625 in. (41.2 mm)
Disk drives	Up to six U.3 Non-Volatile Memory Express (NVMe) solid-state disk (SSD) drives per storage module. Available as 1.6 TB, 3.2 TB, and 6.4 TB drives.
I/O Module	
Physical dimen-	Maximum weight: 5 lb (2.3 kg)

sions	Length: 11.25 in. (285.8 mm) Width: 7.125 in. (181 mm) Height: 3.25 in. (82.6 mm)
Ports	Per I/O module: • Two 10 GbE ports • One 1 GbE (Management) port
PCIe adapters	See https://www.intel.com/ for detailed information about these PCIe adapters: • Intel X550-T2 (10GbE, 2P, Copper) • Intel X710-DA2 (10GbE, 2P, Fibre) • Intel XXV710-DA2 (25 GbE, 2P, Fibre) • Intel E810-CQDA2 (100 GbE, 2P, Fibre) See https://www.marvell.com/ for detailed information about this Fibre-Channel PCIe adapter: • Marvell$^®$ Qlogic$^®$ QLE2770-SR-SP
Control-Panel (Right) Ear	
Port	One USB 2.0 port per system
PSU	
Physical dimen-sions	Weight: 3 lb (1.4 kg) Length: 12.25 in. (311.2 mm) Width: 1.5 in. (38.1 mm) Height: 2.875 in. (73 mm)
Rail Kit	
Weight, in box	9 lb (4 kg)
USB Keyboard and USB Mouse	
Physical dimen-	Weight, in box: 2 lb (.9 kg)

sions	Depth: 8 in. (20.3 cm)
	Width: 19 in. (48.3 cm)
	Height: 2.5 in. (6.4 cm; 2U)
U.S. Power Cords (see "Power Cords" on page 43 for detailed power-cord information)	
Weight, in box	2.6 lb (1.2 kg)
AC Power	
Input power (watts)	31x0: 1247.4W at 100V or 1425W at 200V
	51x0: 1247.4W at 100V or 2970W at 200V
	71x0: 2970W at 200V
	9110: 2970W at 200V
	Notice: These input power ratings are established by National Recognized Testing Laboratories (NRTL) to ensure that the product design and branch circuit supplying power can safely support operation of the product. These tests include ranges of stressful environmental operation parameters, application loading, and optional product configurations. These ratings may not reflect your specific product configuration and operational situation (that is, your system's input power may be less than these ratings), but following these ratings provides a safe installation as tested by the NRTL.
Nominal input voltage	31x0: 100-127V or 200-240V, 50/60 Hz
	51x0: 100-127V or 200-240V, 50/60 Hz
	71x0: 200-240V, 50/60 Hz
	9110: 200-240V, 50/60 Hz
Power factor	0.99
Apparent Power (VA)	31x0: 1260 VA 100V (12.6A x 2) or 1440 VA 200V (7.2A x 2)
	51x0: 1260VA 100V (12.6A x 2) or 3000 VA 200V (15A x 2)
	71x0: 3000 VA 200V (15A x 2)

	9110: 3000 VA 200V (15A x 2)
Protective earth/ground current	0.8 mA maximum for each AC power cord
Thermal load	31x0: 4097 BTUs/hr 51x0: 5868 BTUs/hr 71x0: 6015 BTUs/hr 9110: 6015 BTUs/hr
Environmental	
Operating temperature	41° F to 95° F (5° C to 35° C) For every 800 ft (243.8 m) above 2000 ft (609.6 m), lower the maximum operating temperature by 0.56° F (1° C).
Storage temperature	-38° F to 140° F (-39° C to 60° C), vented
Operating altitude	0 ft to 10,000 ft (0m to 3,048m). In China markets, regulations may limit installations to a maximum altitude of 6562 ft (2000m).
Maximum rate of temperature change during operation	54° F (12° C) per hour
Relative humidity during operation	10% to 80% (noncondensing)
Relative humidity during storage	5% to 95% (noncondensing)
Maximum rate of change of temperature	68°F (20°C) per hour

during storage	
Shock	Operating: Minimum - Fixed setting, Half-Sine wave, 2G, 10ms, 10times Non-operating: Minimum - Fixed in a rack hitting and dropping at a 15mm edge at 300mm/s x3
Vibration	Operating: Minimum - Fixed setting, log sweep, 5 - 30Hz sine wave, 0.25G, 1.5min x2 (3min) all 3 axis Non-operating: Minimum - Fixed setting, 3 - 200Hz random wave, 0.59G, 30min, Z-Axis
Physical Dimensions	
Weight	31x0: 67 lb (30.4 kg) to 75 lb (34 kg), without packaging or pallet 101 lb (45.8 kg), maximum weight with packaging but without pallet 51x0 and 71x0: 71 lb (32.2 kg) to 78 lb (35.3 kg), without packaging or pallet 104 lb (47.1 kg), maximum weight with packaging but without pallet 9110: 71 lb (32.2 kg) to 78 lb (35.3 kg), without packaging or pallet 104 lb (47.1 kg), maximum weight with packaging but without pallet
Height	3.429 in. (87.10 mm)
Width	17.559 in. (446 mm)
Depth	35.826 in. (910 mm)
Acoustics (22C Room Temperature)	
Decibels	31x0: - During pre-boot (first 20 seconds after system is powered on): <=77 dBa (PSU fans run at full speed) - During boot: <= 68.5 dBa

	- During operation: <= 56.1 dBa 51x0 and 71x0: - During pre-boot (first 20 seconds after system is powered on): <=77 dBa (PSU fans run at full speed) - During boot: <= 68.8 dBa - During operation: <= 59.4 dBa 9110: - During pre-boot (first 20 seconds after system is powered on): <=77 dBa (PSU fans run at full speed) - During boot: <= 68.8 dBa - During operation: <= 59.4 dBa
Internal Battery	
Lithium battery	CR2032

Related Topics

- "Safety Precautions" on page 51

- "Memory in ztC Endurance Systems" on page 22

- "Overview of Stratus ztC Endurance Systems" on page 17

Getting Help

The following resources provide help related to your ztC Endurance system:

- View online documentation at the StrataDOC (ztC Endurance Version) page:

    https://ztcendurancedoc.stratus.com

- After you log in to the **Penguin Solutions Customer Service Portal** at https://service.stratus.com, you can:

    - Download software.

    - Search the **Knowledge Base**. The **Knowledge Base** contains technical articles about all Stratus products, including ztC Endurance. In some cases, the online Help directly references these Knowledge Base articles (for example, KB*nnnnnnn*).

> **Notice**: Access to the **Penguin Solutions Customer Service Portal** requires one or more Stratus products to be covered by an active service agreement. If you have questions about your service agreement, contact your authorized Penguin Solutions service representative.

To log in to the Penguin Solutions Customer Service Portal

1. Log in to the **Penguin Solutions Customer Service Portal** at https://service.stratus.com.

   If needed, create a new account as follows:

   a. Click **Register**.

   b. Enter your contact information, including your company email address and registration code, and click **Submit**.

      Your company email address must include a domain name (for example, stratus.com) for a company that is a registered customer of Stratus. The portal sends an email to administrators of the company's account to approve the request.

   c. Upon approval, click the link in the email that you receive from Stratus.

   d. Enter a new password and finish configuring your account.

   If you need assistance creating an account, contact your authorized Penguin Solutions service representative.

To download software from the Penguin Solutions Customer Service Portal

1. Log in to the **Penguin Solutions Customer Service Portal** at https://service.stratus.com.

2. On the **Welcome to Customer Service** page, click **Software Downloads**.

3. On the **Software Downloads** page, select your product, operating system, and model, and then follow the prompts to download the desired software.

To access the Knowledge Base from the Penguin Solutions Customer Service Portal

1. Log in to the **Penguin Solutions Customer Service Portal** at
   https://service.stratus.com.

2. On the **Welcome to Customer Service** page, do one of the following:

   - In the **Search** box, enter keywords or the KB article number (KB*nnnnnnn*) asso-
     ciated with the information you need, and then click the search button.

   - Click **Knowledge**, click the name of a product, and then browse available art-
     icles.

If you cannot resolve your questions with these online resources, and the system is covered by a
service agreement, contact your authorized Penguin Solutions service representative. For information,
see the ztC Endurance Support page at:

   https://www.penguinsolutions.com/en-us/support

Related Topics

-

Using the BIOS Setup Utility

You use the BIOS setup utility to modify BIOS settings and control system hardware.

To start the utility, as the system starts to boot and the Stratus ztC Endurance screen appears, press
either the Delete key or the Esc key until the BIOS setup utility appears.

> **Notice**: Do not change any BIOS settings from their default settings without first being asked
> to do so by Penguin Solutions Customer Service.

Important Considerations When Using the BIOS Setup Utility

Consider the following important issues when you use the BIOS setup utility:

- While you are making changes to the BIOS, make sure that only one compute module is
  running; the other compute module must be powered down. The safest way to do this is to power
  on one compute module at a time. See "Shutting Down and Restarting a Compute Module" on
  page 237.

- Stratus software does change some BIOS settings from their default values, but only at specific
  times. For example, Stratus software enables the OS Watchdog Timer after the software
  installation is complete.

Related Topics

- "Upgrading the System Software" on page 188

REST API

ztC Endurance systems support the Representational State Transfer (REST) application programming interface (API). With access to the REST API operations, you can collect system information and execute some administrative tasks. You can use the REST API to monitor and manage the ztC Endurance system from any location that can access the system over the network.

The system provides the REST API and interactive REST API documentation through the OpenAPI (also known as Swagger™) Specification. You can access the documentation directly on the ztC Endurance system.

To begin exploring the REST API on ztC Endurance systems, see the following sections:

- "Accessing the REST API Documentation" on page 298

- "Logging In for REST API Testing" on page 298

Accessing the REST API Documentation

In a Web browser, open the following URL, which includes one of the hostnames or IP addresses of your ztC Endurance console:

> `https://console_hostname_or_IP address/apidoc`

For example, if the hostname of your ztC Endurance console is `https://sys22.yourcompany.com`, the URL for the REST API documentation is `https://sys22.yourcompany.com/apidoc`.

The URL takes you to the REST API documentation page. This page describes the API operations and syntax, and provides a helpful interface you can use to try the operations on your own system.

Logging In for REST API Testing

> **Notes**:
>
> 1. The system has a limit of 9 login sessions.
>
> 2. Most operations on the REST API documentation page do not require you to perform the following procedure and instead simply require you to enter a username and password after you click **Execute**. However, if you want to use all operations listed on the REST API documentation page, you must perform the following procedure.

1. Open the REST API documentation page.

2. Click **session_api** to expand it. The system displays the syntax for the **session_api** operation.

3. Click **POST /sessions** to expand it, and then click **Try it out**.

4. Under **LoginRequest**, replace `string` next to `username` and `password` with the username and password for a user account with the appropriate level of privileges for the REST API operations that you want to perform. The new username and password appear in bold in the following example.

```
{
   "username": "admin",
   "password": "admin"
}
```

5. Click **Execute**.

6. Scroll down to **Server response** and view **Response body**. If the operation is successful, the response contains a REST API session authentication token (`authToken`). The authentication token appears in bold in the following example. If the operation is unsuccessful, the response might contain a message similar to **Unauthorized**.

```
{
   "authToken": "a8ab7d17f4f3721494eb646ce81dfc24",
   "email": "jbrown@abccompany.com",
   "maxAge": 600000,
   "role": "admin",
   "status": "success"
}
```

7. Upon a successful login, copy the authorization token text located between the quotation marks for use in a subsequent step.

8. At the top of the REST API documentation page, on the right side, click **Authorize**.

9. Under **key (apiKey)**, in the **Value** field, paste the authorization token text, and then click **Authorize**. This saves your authentication token in the request header for any operations that you "execute" from the REST API documentation page.

> **Note**: The token expires after 10 minutes, and you need to repeat this step to create a new token.

To clear the token, use the **Logout** button next to it on the **Authorize** dialog box.

10. Click **Close** or **X** to close the dialog box and return to the REST API documentation page.

11. Proceed with your REST API testing.

End-User License Agreements and Warranty

When you install a ztC Endurance system running Automated Uptime Layer with Smart Exchange, you must read and accept end-user license agreements (EULAs) associated with the system and software. You can read the following EULAs, which are required for Automated Uptime Layer with Smart Exchange, on the Stratus web site:

Stratus Technologies ztC Endurance Series Software End User License Agreement
ztC Endurance Service Terms
Stratus Technologies ztC Endurance Products Limited Return-to-Factory Hardware Warranty

> https://www.stratus.com/services-support/customer-support/platform-support/warranty-info/

You can read the Red Hat Enterprise Linux (RHEL) software license agreement on their EULAs web site (https://www.redhat.com/en/about/eulas). Specifically, read the following:

> https://www.redhat.com/licenses/Red_Hat_GPLv2-Based_EULA_20191118.pdf

Additional EULAs cover third-party software components that are also included with the ztC Endurance system. By accepting the Stratus Technologies ztC Endurance Software EULA and using the ztC Endurance system, you also agree to the terms of the following EULAs:

GNU General Public License (GPL)

> https://www.gnu.org/licenses/gpl-3.0.en.html

GNU Lesser General Public License

> https://www.gnu.org/licenses/lgpl-3.0.en.html

Related Topics

- "Starting a Factory-Installed System for the First Time" on page 161
- "Performing a Full System Software Installation" on page 166

Glossary

A

ActiveService Network (ASN)

The infrastructure that enables communication between a server and the Customer Assistance Center (CAC) or your authorized service representative. Customers can connect to ASN through a modem or over the Internet. Through ASN, alerts can be sent to the CAC or your authorized service representative when an unusual event occurs on the system. ASN also enables the CAC or your authorized service representative to access the server.

B

Baseboard Management Controller (BMC)

Controller used for out-of-band monitoring and management of the ztC Endurance system. Each ztC Endurance system includes two BMCs; the primary BMC monitors the status of system components, even when system power is off. BMCs are equipped with an Ethernet port that can be connected to the Stratus ActiveService Network (ASN). BMCs can communicate with the ASN even when the host operating system is unavailable. This allows remote management, and, if necessary, troubleshooting by the CAC or your authorized Stratus service representative over the internet.

BIOS (basic input/output system)

The program a computer's microprocessor uses to start the computer system after it is powered on. It also manages data flow between the computer's operating system (OS) and attached devices, such as the hard disk, video adapter, keyboard, mouse, and printer.

C

compute module

A customer-replaceable unit (CRU) located at the front of the ztC Endurance system. It is equipped with processors, RDIMMs, BIOS, BMC, fans, and other technology to boot and run the ztC Endurance system. Each system contains two compute modules.

control-panel ear

A module located at the front right of the ztC Endurance system. It houses a USB 2.0 port, power button for the ztC Endurance system, and LEDs that indicate system power status (PWR), fault conditions (ATTN), and module identification status (UID).

customer-replaceable unit (CRU)

Modules that allow on-site personnel to easily remove and replace faulty components with minimal training or tools. ztC Endurance CRU modules are compute modules and their RDIMMs, storage modules and their disk drives, I/O modules and their PCIe adapters, and PSUs.

E

e-Alert

Email alert that the ztC Endurance system sends to system administrators whenever the system detects an event requiring administrator attention.

F

fully qualified domain name

The complete domain name of a specific computer, or host, online. A fully qualified domain name (FQDN) is composed of a hostname and a domain name, separated by a period. An FQDN designates the specific location of an object within the Domain Name System (DNS) hierarchy and allows each entity connected to the Internet to be uniquely identified and located within the Internet framework.

I

I/O module

CRU modules located at the rear of the ztC Endurance system. Each I/O module houses one 1G Ethernet port for connection to a management LAN, two 10G ports for connection to business networks, status LEDs, and five PCIe 4.0 slots for half-height PCIe adapters. The ztC Endurance system contains two I/O modules.

Identify button

Button that illuminates the ID LED on a specific system or component for 30 seconds, so that you can locate it. Identify buttons are found on the front of compute modules, and on the Hardware page of the ztC Endurance console.

K

kickstart file

A text file that contains configuration information for a Red Hat Enterprise Linux (RHEL) installation. The system reads this file at boot time and uses its contents to configure the installation.

L

LUN

A Stratus software-implemented disk device. LUNs are typically duplex, meaning a mirrored set of two raw disks. If a raw disk is missing or not yet configured, the LUN is a single, raw, simplex disk.

M

maintenance mode

A ztC Endurance compute module state that suspends module operation so that it can be serviced. When an active compute module enters maintenance mode, the system performs a Smart Exchange, then activates the standby compute module. Compute modules must be in maintenance mode before they are shut down.

management LAN

The Local Area Network (LAN) used for management of the ztC Endurance system. The I/O modules house a 1G MGMT port for connection to an out-of-band management LAN.

MTBF

The MTBF, or mean time between failure, is the average time that a device works without an error. Regarding the MTBF: 1) Fault and event times are recorded for compute, I/O, and storage modules, as well as PSUs and disks. 2) Events are not faults, but sufficient numbers of certain events, with average in-service time between them falling below a particular threshold, will be counted as a fault. 3) Sufficient numbers of faults, with average in-service time between them falling below a particular threshold, will cause a device to be removed from service or flagged for replacement.

N

Non-maskable interrupt (NMI)

A process, initiated when the system is not responding, where the system controller tries to restart the system by saving the contents of memory to a dump file, and then restarting the operating system. The interrupts are intended to keep all CPUs synchronized while the system controller attempts to resolve the problem.

NVMe (Non-Volatile Memory Express) drive

A protocol designed to use the PCI Express (PCIe) bus to connect SSD (solid-state drive) storage to servers or CPUs. NVMe was created by a consortium of large IT providers in 2008 to provide improved speed and performance.

P

PCIe slot/PCIe adapter

Peripheral Component Interconnect Express (PCIe or PCI-E) is a serial expansion bus standard for connecting a computer to one or more peripheral devices. PCIe provides lower latency and higher data transfer rates than parallel busses such as PCI and PCI-X.

Penguin Solutions Customer Service Portal

Location for downloading ztC Endurance system software, reading knowledge base articles, and getting support. Access to the Penguin Solutions Customer Service Portal requires one or more Stratus products to be covered by an active service agreement, and the creation of an account. The Penguin Solutions Customer Service Portal is located at https://service.stratus.com.

power supply unit (PSU)

Unit that connects the ztC Endurance system to electrical power. Each system includes two PSUs, located at the rear.

R

RDIMM

Registered dual in-line memory module. RDIMMs have a register between the system's memory controller and the dynamic random-access memory modules, thus allowing for greater capacity and increased reliability.

S

secure erase

A process that permanently removes all data from a disk or solid state drive (SSD). After a device has been securely erased, you can safely use it for other purposes, such as backup images.

Simple Network Management Protocol (SNMP)

Network protocol used to monitor the status of network-connected devices. Configuring SNMP settings on ztC Endurance allows remote monitoring of the system by management applications, and allows the system to send messages to designated recipients if alert conditions occur.

Smart Exchange

A proprietary process that maintains system availability and protects against data loss. If the ztC Endurance system detects a degradation or failure within the active compute module, Smart Exchange: 1) Automatically activates the standby compute module and copies over in-process data. 2) Reroutes all I/O and storage connected to the original active compute module to the newly activated compute module. Smart Exchange occurs while the system is running and, except for a very brief pause (usually less than one second), is undetectable to the operating system and customer applications.

standby power

Low-level power that remains present even after the ztC Endurance system has been powered down. Standby power enables the primary Baseboard Management Controller (BMC) to monitor the status of system components, even when system power is off. Using standby power, administrators or ztC Endurance Customer Service (if authorized) can connect to the system through the 1G management port to diagnose problems, or to turn system power on and off from a remote location. To completely shut off power, the power cords must be unplugged.

storage module

CRU modules located at the front of the ztC Endurance system. Storage modules are equipped with status LEDs and include six slots for data disk drives. Each ztC Endurance system contains two storage modules.

Z

ztC Endurance console

A browser-based interface that provides management and monitoring of a ztC Endurance system from a remote management computer. You can perform many administrative

operations from the console because it provides access to the system as a whole, as well as to individual modules, components of modules, and other resources.