ztC Edge User's Guide

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 STRATUS TECHNOLOGIES, STRATUS MAKES NO WARRANTY OR REPRESENTATION OF ANY KIND WITH RESPECT TO THE INFORMATION CONTAINED HEREIN, INCLUDING WARRANTY OF MERCHANTABILITY AND FITNESS FOR A PURPOSE.

Stratus Technologies 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 Stratus documents (a) is the property of Stratus Technologies Ireland, Ltd. 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.

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

Copyrights

Manual Name: *ztC Edge User's Guide*

Product Release Number: Stratus Redundant Linux Release 1.3.0.0

Publication Date: Thursday, April 29, 2021

Stratus Technologies, Inc.

5 Mill and Main Place, Suite 500

Maynard, Massachusetts 01754-2660

Table of Contents

Part 1: ztC Edge User's Guide **1**

Chapter 1: Introduction to ztC Edge Systems **1**

ztC Edge System Overview 1

ztC Edge System Description 2

Physical Machines and Virtual Machines 2

Administrative Operations 3

Alerts 3

Remote Support 4

Lights-Out Management 4

Third-party Management Tools 5

Modes of Operation 5

High Availability Operation 6

ALSR Configurations 6

ALSR and Quorum Service 6

Quorum Servers 7

Network Architecture 8

A-Link and Private Networks 8

Business and Management Networks 9

Network Segmentation Fault Detection and Remediation 10

System Usage Restrictions 10

QEMU 10

Accessing the Host Operating System 11

Chapter 2: Getting Started **13**

Planning 13

Safety Precautions 14

System Requirements Overview 14

System Hardware 14

RAM 14

Disk Space 14

Network 15

IP Addresses 15

Ports 15

General Network Requirements and Configurations 16

Business and Management Network Requirements 16
A-Link and Private Network Requirements 18
ztC Console Requirements 18
Compatible Internet Browsers 18
Power Requirements and Considerations 19
Deployment 19
Connecting Power 19
UPS (Optional) 20
Deploying the System 21
Connecting Ethernet Cables 23
Mapping Your Keyboard 25
To configure your keyboard layout after deployment: 25
Recording the Management IP Address 26
Post-Deployment Tasks 26
Obtaining System IP Information 27
Logging On to the ztC Console for the First Time 27
Registering the System and Acquiring a Permanent License 29
Redeploying a ztC Edge System 33
Connecting a Second Business Network 34
Chapter 3: Using the ztC Console 35
The ztC Console 36
Logging On to the ztC Console 37
The Dashboard Page 37
Resolving Outstanding Alerts on the Dashboard 38
The System Page 39
Powering On the System 39
Rebooting the System 40
Shutting Down the System 41
The Preferences Page 43
Specifying Owner Information 45
Managing the Product License 45
Configuring IP Settings 48
Configuring Quorum Servers 50
Configuring Date and Time 52
Configuring Secure Connections 53

Configuring VM Devices 57
Configuring e-Alerts 58
Configuring SNMP Settings 60
Configuring Remote Support Settings 62
Configuring Internet Proxy Settings 64
The Alerts History Page 65
The Audit Logs Page 65
The Support Logs Page 66
Creating a Diagnostic File 66
Uploading a Diagnostic File to Customer Support 67
Deleting a Diagnostic File 68
The Physical Machines Page 69
Physical Machine Actions 70
Physical Machine States and Activities 71
The Virtual Machines Page 72
Virtual Machine Actions 73
Virtual Machine States and Activities 76
The Volumes Page 77
The Networks Page 79
Setting the MTU 80
The Virtual CDs Page 80
The Upgrade Kits Page 81
Creating a USB Medium with System Software 81
The Users & Groups Page 82
Managing Local User Accounts 83
Chapter 4: Upgrading Stratus Redundant Linux Software 85
Upgrading Stratus Redundant Linux Software Using an Upgrade Kit 85
Chapter 5: Managing Physical Machines 89
Maintenance Mode 89
Powering On a Physical Machine 91
Rebooting a Physical Machine 91
Shutting Down a Physical Machine 92
Load Balancing 94
Modes of Operation 94
Troubleshooting Physical Machines 95

Recovering a Failed Physical Machine (Manual) 96
Chapter 6: Managing Virtual Machines 101
Planning Virtual Machine Resources 102
Planning Virtual Machine vCPUs 102
Planning Virtual Machine Memory 104
Planning Virtual Machine Storage 105
Planning Virtual Machine Networks 106
Creating and Migrating Virtual Machines 106
Creating a New Virtual Machine 107
Copying a Virtual Machine 110
Migrating a Physical Machine or Virtual Machine to a System 113
Importing an OVF File 124
Replacing/Restoring a Virtual Machine from an OVF File 130
Exporting a Virtual Machine 134
Mounting a USB Device or Network-mounted Folder on the ztC Edge System 138
Managing Windows Drive Labels 140
Configuring Windows-based Virtual Machines 141
Updating the VirtIO Drivers (Windows-based VMs) 142
Creating and Initializing a Disk (Windows-based VMs) 145
Installing Applications (Windows-based VMs) 145
Configuring Linux-based Virtual Machines 146
Creating and Initializing a Disk (Linux-based VMs) 146
Installing Applications (Linux-based VMs) 147
Managing the Operation of a Virtual Machine 147
Starting a Virtual Machine 148
Shutting Down a Virtual Machine 149
Powering Off a Virtual Machine 150
Opening a Virtual Machine Console Session 151
Renaming a Virtual Machine 153
Removing a Virtual Machine 153
Managing Virtual Machine Resources 154
Reprovisioning Virtual Machine Resources 155
Creating a Volume in a Virtual Machine 157
Attaching a Volume to a Virtual Machine 158
Detaching a Volume from a Virtual Machine 159

Removing a Volume from a Virtual Machine 161
Renaming a Volume on the ztC Edge System 162
Expanding a Volume on the ztC Edge System 163
Recovering Virtual Machine Resources 164
Managing Virtual CDs 165
Creating a Virtual CD 165
Inserting a Virtual CD 167
Ejecting a Virtual CD 167
Booting from a Virtual CD 168
Renaming a Virtual CD 169
Removing a Virtual CD 169
Advanced Topics (Virtual Machines) 170
Assigning a Specific MAC Address to a Virtual Machine 170
Selecting a Preferred PM for a Virtual Machine 172
Forcing a VM to Boot 172
Configuring the Boot Sequence for Virtual Machines 176
Resetting MTBF for a Failed Virtual Machine 177
Attaching a USB Device to a Virtual Machine 178
Chapter 7: Maintaining Physical Machines 181
Replacing Physical Machines (Automated) 181
Replacing Physical Machines (Manual) 183
Part 2: Supporting Documents 189
Chapter 8: ztC Edge Release 1.3.0.0 Release Notes 190
Accessing Stratus Knowledge Base Articles 190
Bug Fixes 191
Fixed in ztC Edge Release 1.3.0.0 191
Fixed in ztC Edge Release 1.2.1.0 191
New Features and Enhancements 191
New in ztC Edge Release 1.3.0.0 192
New in ztC Edge Release 1.2.1.0 192
New in ztC Edge Release 1.2.0.5 192
Important Considerations 192
Upgrading to ztC Edge Release 1.3.0.0 192
Using Intel Active Management Technology (AMT) for Lights-Out Support 193
Redeploying a ztC Edge System 193

Deploying ztC Edge Nodes at Separate Physical Sites 194
Other Important Considerations 194
Known Issues 194
Do Not Click the Remove Button on the Physical Machines Page 194
NFS or SMB/CIFS Mounts May Delay Failover 194
Importing or Exporting a VM Through a Windows Share (CIFS/SMB) May Fail 194
VMs Reported as Broken Instead of Degraded When A-Link Is Offline 194
Ejected VCD Still Displayed in a Linux-based VM Console 195
Cannot Re-mount USB Media After Ejecting from Windows File Explorer in a VM 195
VM Network Remains Offline After Using P2V/V2V to Migrate from a VMware Source 195
VM Network Remains Offline After Importing a Windows-based VM 195
Ensure that Both Nodes Are Connected to the Network Before Creating a VM 195
Ensure that Both Nodes Are Running Before Removing a VM 196
Invalid VM Names Prevent Some VM Operations from Executing 196
Some Browsers Unable to Connect a VNC When Using https 196
Cannot Mount USB Media with exFAT File System 197
Reboot Required when Changing Node IP Address or Netmask Network Settings 197
Changing the MTU of Business Networks Causes VM Migration and Possibly Node Failover 197
Getting Help 197
Chapter 9: System Reference Information 198
Compatible Guest Operating Systems 198
Important Physical Machine and Virtual Machine Considerations 199
Virtual Machine Recommendations and Limits 199
Recommended Number of CPU Cores 200
Important Considerations 200
Creating an ALSR Configuration 200
Creating the Configuration 205
A Typical ztC Edge System 205
An ALSR Configuration With a Quorum Server 206
ALSR VLAN Requirements 207
From Initial Deployment to Completing the ALSR Configuration 207
Meeting Network Requirements 209
Locating and Creating the Quorum Server 210
Locating the Quorum Computer 211
Adding an Alternate Quorum Computer 211

Quorum Computer Requirements 212
Downloading and Installing the Quorum Service Software 212
Completing the Configuration 213
Configuring the Quorum Service Port 213
Verifying the Quorum Service Port 214
Configuring the Quorum Server Within the ztC Console 214
Verify the Configuration and (Re-)Join VMs 215
Understanding Quorum's Effect on System Behavior 216
Example 1: A System Without a Quorum Server Experiences a Split-brain Condition 216
A Catastrophic Fault 217
Fault Handling 217
Recovery and Repair 217
Example 2: An ALSR System With a Quorum Server Avoids a Split-brain Condition 218
A Catastrophic Fault 219
Fault Handling 219
Recovery and Repair 220
Example 2, Modified: The Quorum Server Is Unreachable During the Catastrophic Fault 220
Example 2, Modified: The Quorum Server Is Unreachable With No Catastrophic Fault 221
Recovering From a Power Failure 221
Accessing Knowledge Base Articles 221
REST API 222
Chapter 10: SNMP 224
Obtaining System Information with snmptable 224

Part 1: ztC Edge User's Guide

The *ztC Edge User's Guide* describes ztC Edge systems, how to deploy them, and how to use them.

For system descriptions including modes of operation and storage and network architecture, see:

- Introduction to ztC Edge Systems

For planning and deployment information, see:

- Getting Started

The following topics describe how to administer ztC Edge systems:

- Using the ztC Console
- Upgrading Stratus Redundant Linux Software
- Managing Physical Machines
- Managing Virtual Machines
- Maintaining Physical Machines

1

Chapter 1: Introduction to ztC Edge Systems

See the following topics for an introduction to ztC Edge systems:

- ztC Edge System Overview
- Modes of Operation
- Network Architecture
- System Usage Restrictions

ztC Edge System Overview

A ztC Edge system with two nodes provides automated recovery with no lost data in the event of a hardware failure.

See the following topics for descriptions of system features and capabilities.

- ztC Edge System Description
- Physical Machines and Virtual Machines
- Administrative Operations
- Alerts
- Remote Support
- Lights-Out Management
- Third-party Management Tools

ztC Edge System Description

Stratus Redundant Linux software allows two individual ztC Edge computers (with the appropriate license) to work as a single, highly-available or fault-tolerant system. Each computer is called a physical machine (PM) or node. Both PMs:

- Run the same host operating system (CentOS)
- Contain replicated virtual machines and storage (synchronized via direct Ethernet links between the two PMs)
- Support virtual machines running supported guest operating systems

For more information about the configuration of PMs in a ztC Edge system, see System Requirements Overview.

Stratus Redundant Linux software can also run on a single PM when the system is licensed for one node. In this configuration, the system is simplex, it is not fault-tolerant or highly-available, and during normal operation, the system displays network faults.

Related Topics

System Requirements Overview

Compatible Guest Operating Systems

Network Architecture

Physical Machines and Virtual Machines

A ztC Edge system, which runs Stratus Redundant Linux software on two physical machines (PMs), also referred to as nodes, transparently protects applications by creating redundant virtual machines (VMs) that run on both nodes.

Stratus Redundant Linux software also runs on a system with one PM, when licensed for one node. For information on systems licensed for one node, see ztC Edge System Description. The information in the rest of this topic applies to systems licensed for two nodes.

The Stratus Redundant Linux management software can create a guest VM from scratch. It can also import existing VMs from other environments and convert them into guest VMs. By creating an identical instance of the selected VM on a second host PM, the management software automatically provides high-availability (HA) protection of the VM. The system administrator manages this single entity from a separate, browser-based management console called the ztC Console.

Neither the application nor the user is exposed to the redundant computing resources on the two host PMs. The application sees only one hostname, one MAC address for each network interface presented to the VM, and one IP address for each VM network interface presented to the VM. A system administrator loads and configures the applications on the guest VM–just as if the system administrator were loading them onto a physical server. If a fault or failure occurs in a disk or network device, the software automatically redirects I/O to the paired host PM. Though redundancy is lost until the failure is repaired, the VM continues to operate normally. The application continues to execute as if nothing had happened. The redundancy, fault detection, isolation, and management are completely transparent to the Windows or Linux environment and the application running within it. Repair of the PM is equally transparent and automatic. When a failed component on the PM is repaired, the software automatically incorporates the repaired components into the protected environment of the guest VM and restores redundancy transparently.

Related Topics

Using the ztC Console

The Physical Machines Page

The Virtual Machines Page

Administrative Operations

You can perform many administrative operations on the ztC Edge system from the ztC Console, a browser-based interface that provides access to the system as a whole as well as to physical machines (PMs), virtual machines (VMs), and other resources. For information, see The ztC Console.

Alerts

ztC Edge system alert messages notify the system administrator whenever an item needs attention. These can include:

- Configuration tasks that should be performed
- Notification of system operational states
- System problems that require attention

Click **Dashboard** in the left-hand navigation panel to see Alert messages and their descriptions. Click **Alerts** in the left-hand navigation panel to see the Alert log.

The following icons indicate the state of an alert message.

 Informational

 Normal or OK state

 Minor, warning, or inconsistent state

 Moderate state

 Broken, failed, or severe state

Remote Support

To access the ztC Edge system's remote support features, click **Preferences** in the left-hand navigation panel. From there, you can configure support and proxy specifications by selecting the following:

- **Support Configuration**–Configure settings to allow remote support access of your system by your authorized Stratus service representative and to enable your system to send health and status notifications to your authorized Stratus service representative. See Configuring Remote Support Settings for details.
- **Proxy Configuration**–Enables you to configure a proxy server for access to the Internet. See Configuring Internet Proxy Settings for details.

Lights-Out Management

ztC Edge systems incorporate Intel® Active Management Technology (AMT) lights-out support, which is disabled by default. You can enable and configure this support by pressing **Ctrl-P** while the BIOS splash screen is displayed during system startup. For important information about AMT configuration and restrictions, access the Knowledge Base to search for the article *AMT and Remote Access in ztC Edge* (KB-8219). See Accessing Knowledge Base Articles.

AMT features are accessible on the **P1** network port of the system.

Third-party Management Tools

You can install third-party management tools on ztC Edge systems. Examples of such tools include vendor-specific management/monitoring utilities, enterprise management/monitoring utilities, and other miscellaneous management/monitoring software. Note the following:

- In general, management tools that run on the host operating system (CentOS) should run on ztC Edge systems. Possible exceptions are tools that manage/monitor the CentOS KVM-based virtualization. To manage/monitor ztC Edge virtualization, use the integrated ztC Edge management tools.
- Before deploying your ztC Edge system, Stratus recommends that you verify that it operates properly with the management tools installed and operational.
- Stratus recommends that you set up a non-root account for third-party management tools.
- You can access your ztC Edge system via the management network using the IP address(es) specified during the installation process (or supplied by the DHCP server if the interface was configured for DHCP during install).
- If you install third-party management tools in the host operating system of a physical machine (PM) and you need to replace the PM in the future, remember to reinstall the tools on the replacement PM.

Note: Third-party management tools have the potential of destabilizing the environment of the host operating system and system software. You may need to remove management tools that consume excessive RAM or disk space, or that are otherwise suspected of destabilizing the product. Follow the recommendation of your authorized Stratus service representative.

For information about accessing the host operating system, see Accessing the Host Operating System.

Related Topics

Getting Started

System Reference Information

Modes of Operation

Stratus Redundant Linux provides High Availability (HA) operation for VMs. HA operation achieves its level of redundancy by using a pair of physical machines (PMs). See High Availability Operation.

Stratus recommends configuring quorum service for HA operation. The quorum service prevents a condition called *split-brain* where both PMs of an HA operation pair are running independently of each other; for information, see Quorum Servers.

High Availability Operation

In HA operation, Stratus Redundant Linux automatically detects, isolates, and handles most hardware faults, thereby keeping your applications running. With HA remote-support technology, the software notifies the Stratus support center of various issues, indicating the type of fault and its exact location. This combination of automatic fault detection, isolation, and remote-support technologies ensures speedy access to expert support technicians and rapid problem resolution.

HA operation offers basic failover and recovery, with some faults requiring an (automatic) VM reboot for recovery, and return to HA operation:

- Eliminates downtime for many, but not all, CPU, memory, I/O, or other physical machine (PM) failures.
- Handles failures without IT intervention.
- Provides continuous, active validation of all components.
- Assures redundancy and recovery at all times.

HA is suitable for applications that can tolerate occasional interruptions of a few minutes.

Related Topics

The Virtual Machines Page

Using the ztC Console

ALSR Configurations

An *automated local site recovery* (ALSR) *configuration* connects two physical machines in two separate sites. It is a disaster-tolerant deployment that maintains hardware redundancy as well as redundancy of physical computer rooms and the buildings containing them. Because of the geographic separation, an ALSR configuration requires careful planning of component placement and more complex networking topologies. **For ALSR configurations, Stratus strongly recommends that you use the quorum service because an ALSR configuration exposes the A-Link networks to other potential failure scenarios.** (ALSR configurations are not available to systems licensed for one node.)

Meeting Network Requirements lists the requirements for networks in a ALSR configuration.

ALSR and Quorum Service

In an ALSR configuration, configure two quorum-service computers in compliance with the best practices recommended for quorum deployment (see Quorum Servers and Locating and Creating the Quorum Server).

In any ALSR configuration, a preferred quorum-service computer is located in a third facility, and an alternate is located in a fourth site (or carefully placed in the third). The networks are interconnected.

Quorum-service computers should be as isolated as possible. If both must be placed in a common (third) site, make sure that they do not depend on common power sources.

Physical connectivity between an ztC Edge PM and the quorum-service computers must not route through the other PM's site.

Placing a quorum-service computer in the same site as one of the ztC Edge PMs ensures data integrity. However, some site failures may then require that the VMs be shut down until manually recovered.

The management network physically connects the ztC Edge PMs and the quorum-service computers. For this to work properly, you must configure each ztC Edge PM to use a different gateway to reach the quorum-service computers. If the two PMs use the same gateway to reach the quorum-service computers, data integrity is ensured during failures. However, some site failures may then require that the VMs be shut down until manually recovered.

Related Topics

Creating an ALSR Configuration

Network Architecture

Quorum Servers

A *quorum service* is a Windows operating system-based service deployed on a server distinct from the two servers (physical machines or PMs) of a dual-node ztC Edge system. Quorum servers provide data integrity assurances and automatic restart capabilities for specific failures in a ztC Edge environment. Stratus strongly recommends using quorum servers, especially for ALSR operation. You can configure the two PMs of a ztC Edge system with 0, 1, or 2 quorum servers.

Quorum servers ensure the integrity of VMs against multiple network failure scenarios, including split-brain, and provide for unattended startup of VMs after specific failures. Quorum server communication occurs via the management network.

Quorum servers are particularly important in ALSR configurations. Best practice for ALSR is to place a preferred quorum computer in a third facility and an alternate quorum computer in a fourth facility. However, you can also place the alternate quorum service computer with the preferred quorum computer and still obtain satisfactory service. See ALSR Configurations for additional information.

If only two sites are available (thereby preventing the best practices configuration described above) and if one PM goes down and the surviving PM is unable to communicate with the quorum server (for example, because

it is on the same site as the down PM), the VMs at the surviving site are automatically shut down to avoid running in split-brain.

Related Topics

Creating an ALSR Configuration, which discusses quorum servers

Configuring Quorum Servers

Network Architecture

Ethernet networks provide pathways for communications of a system. The main types of Ethernet networks are:

- *Availability Link networks*, or *A-Link networks* (on the blue (**A2** or •) and yellow (**A1** or • •) network ports) on ztC Edge systems licensed for two nodes are assigned to virtual machines (VMs) and are used to synchronize data or migrate VMs between two PMs. One A-Link network (on the blue (**A2** or •) network port), is a *private network* (priv0) that connects the two ztC Edge PMs. For more information, see A-Link and Private Networks. (Systems licensed for one node do not provide A-Link network functionality.)
- *Business networks* (on the **P1** network port, and **P2** if enabled) on all ztC Edge systems allow your applications to connect to your existing network. One business network (on the **P1** network port) is also a *management network* (ibiz0, sometimes referred to as network0) that connects to the ztC Console and is used by the quorum servers. For more information, see Business and Management Networks.

A ztC Edge system also provides a network segmentation detection mechanism. For information, see Network Segmentation Fault Detection and Remediation.

A-Link and Private Networks

Every ztC Edge system licensed for two physical machines (PMs, which are also referred to as nodes) requires a network for private management traffic between the two PMs. This private network is referred to as *priv0*, which is a physical, direct Ethernet, or VLANed connection between the nodes. Priv0 is used for peer node discovery and can have no other entities on it that respond to IPv4 broadcasts.

In addition to priv0, each system licensed for two nodes has A-Link networks to increase data-replication performance between the PMs. A-Link networks let you sync disks, shunt networks, migrate VMs, perform heartbeat checks, and sync fault-tolerant memory.

The A-Links and priv0 are connected between the PMs in the same manner. The A-Links are connected between the blue (•) and yellow (• •) network ports of each PM, where priv0 is shared with the A-Link on the blue (•) network.

The simplest priv0 consists of a single Ethernet cable (crossover or straight-through) that directly connects an embedded Ethernet port on each server. If a networking device other than a single Ethernet cable is used for priv0, see ALSR Configurations.

Related Topics

Business and Management Networks

A-Link and Private Network Requirements

Network Architecture

Business and Management Networks

All Ethernet ports—other than those used by A-Link networks and the private network—are considered business-network ports. Guest operating systems use business-network ports to connect to your network.

One business network is the *management* network, and each PM has a single management network that is referred to as *ibiz0* and uses the network labeled **P1**. The management network accesses the ztC Console and handles miscellaneous management tasks and the quorum server. These management tasks include:

- Sending call-home messages and e-alerts
- Checking the status of licenses
- Each PM's communication with the ztC Console
- Failover function of priv0 (for systems with dual nodes)
- Communication between the two nodes (for systems with dual nodes)
- Communication with the quorum server (if one exists)

You set up the management network when you deploy the system. You can also set up business networks for any business-network ports that are physically connected during deployment. To connect a second business network after the deployment is complete, see Connecting a Second Business Network.

Related Topics

A-Link and Private Networks

Business and Management Network Requirements

Network Architecture

Network Segmentation Fault Detection and Remediation

A network fault that occurs such that the two ends of a shared network cannot communicate with each other, but each side still has external network connectivity, is referred to as a *network segmentation fault*.

ztC Edge system provides a *network segmentation detection mechanism* that places the active VM on the node that has the most external network connectivity when the system detects this fault. As part of this feature, the ztC Edge system constantly sends UDP packets over the business network interface between the active node and the stand-by node. The system's network segmentation logic detects a fault when this packet flow is interrupted while both sides still have an active network link. In this fault scenario, both nodes still have active network connections, so the fault lies in a switch that is external to the ztC Edge system.

When this case is detected, the ztC Edge system handles the fault based on logic that determines which side has better external connectivity. The ztC Edge system makes this fault-handling decision by continually monitoring incoming broadcast/multicast traffic to determine which node has the most incoming traffic. In this fault case, if the VM is not already active on the node with the most incoming network traffic, the ztC Edge system fails the VM network over to this node. The fault detection feature requires no user configuration since it is basing the decision on traffic that normally occurs on any system.

Related Topics

Network Architecture

System Usage Restrictions

Observe the restrictions to system usage that are described in the following topics:

- QEMU
- Accessing the Host Operating System

QEMU

ztC Edge systems support the open-sourced hypervisor QEMU ("Quick EMUlator"), which performs hardware virtualization. When used as a virtualizer, QEMU executes the guest code directly on the host CPU, achieving a high level of performance.

ztC Edge users should make no changes to the QEMU virtualization engine or to its configuration.

Accessing the Host Operating System

After you complete the ztC Edge deployment, you can access the host operating system (CentOS) locally at the PM's physical console, or you can access it remotely by using a secure shell (SSH) client.

To log on to the host operating system with an SSH client, use the management IP address specified during deployment (or supplied by the DHCP server, if the interface was configured for DHCP during deployment). If needed, you can locate the management IP address for a PM as described in this topic.

Caution: Do not update the CentOS host operating system of the ztC Edge system from any source other than Stratus. Use only the CentOS release that is installed with the Stratus Redundant Linux software.

Note: To ensure that administrative commands will work properly, log on to the physical console or IP address of the primary PM (unless you specifically need to operate on components in the secondary PM). Do not connect to the system IP address, as it can move from PM to PM.

The default password for the root account is **KeepRunning**.

Note: For security reasons, change the username and password as soon as possible.

For information about using third-party management tools on CentOS, see Third-party Management Tools.

To locate the IP address of each PM in the ztC Console

1. Click **Preferences** in the left-hand navigation panel to open the **Preferences** page.
2. Under **System**, click **IP Configuration**.
3. Record the **IP address** of each PM, **node0** and **node1**.
4. Click **Physical Machines** in the left-hand navigation panel to open the **Physical Machines** page.
5. Record which PM is the primary node for the system, displayed as **node*n* (primary)**. In most cases, log on to the IP address of the primary node to ensure that administrative commands will work properly.

To access the host operating system from a Windows-based system

You can download and use PuTTY, a suite of open-source SSH clients:

http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html

In particular, the `putty.exe` client allows you access a shell to execute programs on the command line of the host operating system. PuTTY also includes the `pscp.exe` command-line utility that allows you to securely transfer files from a remote system to the host operating system.

If you prefer a secure copy (SCP) client with a graphical user interface, you can also try the open-source WinSCP utility:

http://winscp.net/eng/index.php

To access the host operating system from a Linux-based system

On many Linux- and UNIX-based systems, SSH utilities are already installed and enabled by default. See ssh(1) and scp(1) for information about how to use these utilities.

2

Chapter 2: Getting Started

The following topics describe the ztC Edge planning, deployment, and post-deployment tasks:

- Planning
- Deployment
- Post-Deployment Tasks

Planning

See the following topics for information about planning your system configuration.

- Safety Precautions
- System Requirements Overview
- General Network Requirements and Configurations
- Business and Management Network Requirements
- A-Link and Private Network Requirements
- ztC Console Requirements
- Compatible Internet Browsers
- Power Requirements and Considerations
- Creating an ALSR Configuration (if applicable to your configuration)

After you have planned the system configuration, continue with Deployment.

Safety Precautions

Before getting started, please read the following important safety precautions.

Warning: Ensure the voltage of the power source is correct before connecting the product.

Warning: Servicing to be performed by qualified service personal, no user-serviceable components.

Warning: Risk of explosion if battery is replaced by an incorrect type. Dispose of used batteries according to the instructions.
IL Y A RISQUE D'EXPLOSION SI LA BATTERIE EST REMPLACEE PER UNE BATTERIE DE TYPE INCORRECT. METTRE AU REBUT LES BATTERIES USAGEES CONFORMEMENT AUX INSTRUCTIONS

Warning: Hot Surface Do Not Touch.

System Requirements Overview

A ztC Edge system can support multiple virtual machines (VMs) and a remote management computer (that is, a general-purpose PC) that can run the ztC Console.

ztC Edge System Hardware specifications and requirements are summarized below.

For information on guest operating systems, see Compatible Guest Operating Systems.

System Hardware

RAM

Each PM has 32 GB of RAM (physical memory).

Disk Space

Each PM contains a 512 GB solid-state drive (SSD), of which approximately 475 GB is available for VMs.

Network

Each PM has four 1-Gb Ethernet ports, including:

- On systems licensed for two nodes:
  - A combined A-link and priv0 private network port, labeled blue (•)
  - A second, dedicated A-link port, labeled yellow (• •)
- On systems licensed for two nodes or one node:
  - A combined business network and management network port, labeled **P1**
  - An optional business network port, labeled **P2**

The system also supports Intel® Active Management Technology (AMT) lights-out support, which you can access over the **P1** port of each PM.

ALSR configurations have different network requirements. For information, see Meeting Network Requirements .

See Network Architecture, A-Link and Private Networks, and Business and Management Networks for more information.

IP Addresses

Each ztC Edge system must have a static IPv4 IP address assigned for use by the management software. Obtain IP addresses for DNS primary and secondary servers, and gateway and subnet mask information for your management network, from your IT network administrator. See Obtaining System IP Information for more information.

Ports

ztC Edge systems use port 443 in the local firewall for HTTPS communications, port 22 for ssh, and 5900-59*nn* for each active VNC associated with each VM. Firewalls must allow traffic through the appropriate ports. Firewalls must permit VMs to contact quorum service computers using UDP port 4557. For additional information on TCP and UDP ports, access the Knowledge Base to search for the article *TCP and UDP ports used by ztC Edge* (KB-9357). See Accessing Knowledge Base Articles.

Related Topics

Important Physical Machine and Virtual Machine Considerations

Virtual Machine Recommendations and Limits

Planning Virtual Machine Resources

Configuring IP Settings

General Network Requirements and Configurations

> **Note**: ALSR networks have some additional and different network requirements and recommendations. See Creating an ALSR Configuration in addition to the information below.

Before you deploy a ztC Edge system, make sure your network meets the following requirement:

- ztC Edge systems utilize full IPv4 and IPv6 protocol access, including IPv6 multicast. Any obstruction of this traffic may prevent a successful deployment or compromise the availability of a running ztC Edge system.

In addition, see the following topics for the requirements specific to each network type:

- A-Link and Private Network Requirements
- Business and Management Network Requirements

Business and Management Network Requirements

Business and management networks have the following requirements:

- The networks use IPv6 link-local addressing.
- The networks support an MTU value of up to 9000.
- The networks do not support bonding or VLAN trunking.
- Virtual machines (VMs) can use IPv4, IPv6, and other Ethernet protocols.
- All business networks can be used for IPv6 host access if your site has SLAAC or DHCPv6 enabled.
- To reach the ztC Console, use ibiz0, which is the IPv4 address that migrates to the primary management physical machine (PM). Each PM also has its own ibiz0 IPv4 address on the management network.
- Each PM requires at least one business network (specifically, the management network).

To ensure that Ethernet traffic flows unobstructed to and from VMs from either PM:

- The switch ports connected to business networks must not filter ARP packets, including gratuitous ARP packets. A ztC Edge system sends gratuitous ARP packets on behalf of guest VMs in order to prompt Ethernet switches to update their port-forwarding tables to direct VM traffic to the appropriate physical Ethernet port on the appropriate PM.
- The switch ports connected to business networks must allow layer2 multicasts (address: 01:E0:09:05:00:02) with ethertype: 0x8807.
- If you configure RHEL or CentOS guests to have multiple NICs on same subnet, you may experience guest network connectivity issues due to asymmetric routing. To avoid this problem, modify the **/etc/sysctl.conf** file on the guest Virtual Machine (VM) to contain the following lines, save the file, and reboot the VM.
  - `net.ipv4.conf.default.rp_filter = 2`
  - `net.ipv4.conf.all.rp_filter = 2`
- Do not issue the `ifdown` command from a PM's host OS to temporarily bring down a VM's business network connection (ibiz*x*). Doing so will disconnect the physical interface from its bridge and cause the VM to become unreachable over the network. Instead, use the `ifconfig down` command.
- The switches connected to business networks must not enable any MAC address security features that would disable the movement of a MAC address from one business link to the matching business link on the other PM.
- For optimal failover response, configure any switches connected to your system to have MAC aging timeout values of less than one second.

If these requirements are not met, or if the switch does not properly update its forwarding table when a VM is migrated from one ztC Edge PM to the other PM, the VM may experience a blackout in which network traffic is not properly directed to and from the VM.

Related Topics

Network Architecture

Business and Management Networks

A-Link and Private Network Requirements

A-Link and private networks have the following requirements:

- The networks use IPv6 link-local addressing.
- All A-Link and private networks on one PM of a ztC Edge system must be in the same L2 broadcast domain as its matching links on the other physical machine (PM), without any protocol filtering.
- Ethernet packets sent between two PMs of a system must not be obstructed or rate-limited. Ensure that they are not routed or switched by any L3 network infrastructure.
- The speed of A-Link networks should be equal to or greater than the speed of business or management networks.
- Network traffic for storage replication between PMs is sent over A-Link networks.
- Private networks have no network hosts connected other than the ztC Edge end-points.

Related Topics

A-Link and Private Networks

ztC Console Requirements

The ztC Console provides browser-based remote management of the ztC Edge system, its physical machines (PMs), and virtual machines (VMs).

- Your computer must be able to access the subnet containing the ztC Edge management network (which is enabled on the network port labeled **P1**).
- Use a supported browser. See Compatible Internet Browsers.

For more information, see Using the ztC Console.

Compatible Internet Browsers

A browser is used to connect to the ztC Console. Use only browsers that are compatible with ztC Edge systems. Using an incompatible browser can result in some rendering problems and the omission of some wizards.

The following browsers are compatible with ztC Edge systems.

Compatible Browsers	Release
Microsoft Internet Explorer™	IE9 or newer[1]
Mozilla® Firefox®	25 or newer
Google® Chrome™	31 or newer

Power Requirements and Considerations

To ensure maximum availability, Stratus strongly recommends that each physical machine (PM), or node, in a ztC Edge system connects to a separate, independent power source.

See Connecting Power for illustrations of some sample power-connection configurations.

Deployment

When you deploy the system for the first time:

> **Note**: If you have already deployed and configured a system, and you need to prepare it for deployment at a new site, see Redeploying a ztC Edge System.

1. Review the network cabling information. If necessary, make changes in your network. See Connecting Ethernet Cables.
2. Deploy the system. See Deploying the System.

When the deployment is complete, see Post-Deployment Tasks.

Related Topics

Upgrading Stratus Redundant Linux Software

Connecting Power

To connect power, configure an ztC Edge system licensed for two nodes with redundant power supplies connected to separate sources. You can optionally use uninterruptible power supplies (UPS), as shown below.

After connecting power, return to Deploying the System.

[1] IE8 is not recommended and it does not support some ztC Edge functionality.

UPS (Optional)

The illustrations show how to connect one or two optional UPS units to an ztC Edge system licensed for two nodes.

> **Note**: Stratus recommends that you use two UPS units connected to separate and independent power sources. With two power sources, the system continues to receive power in the event that one power source fails.

Single UPS:

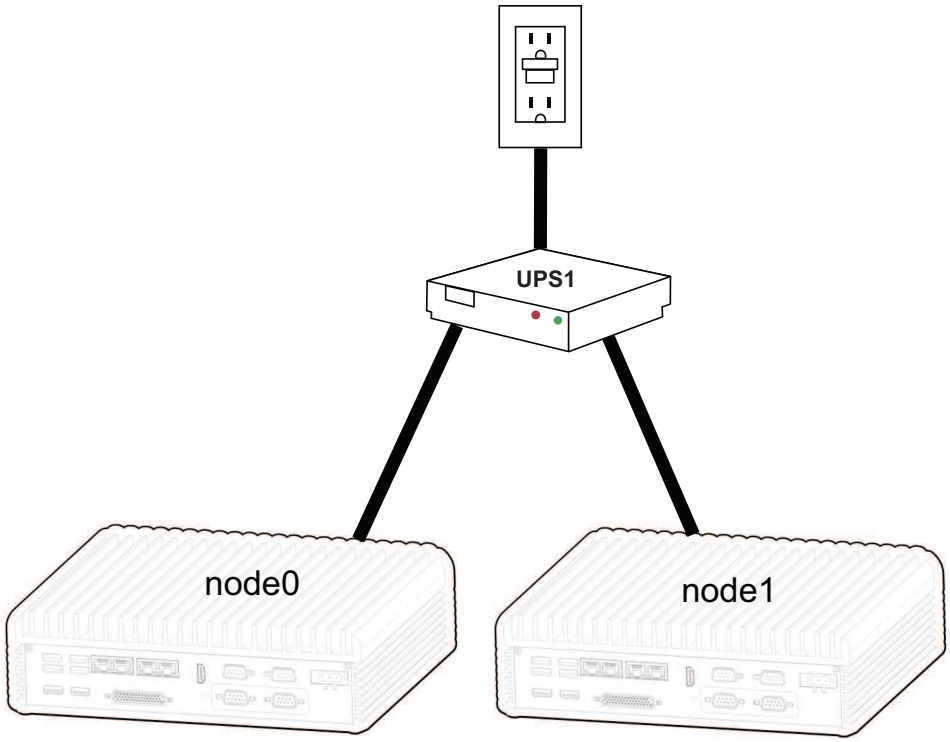

Dual UPS:

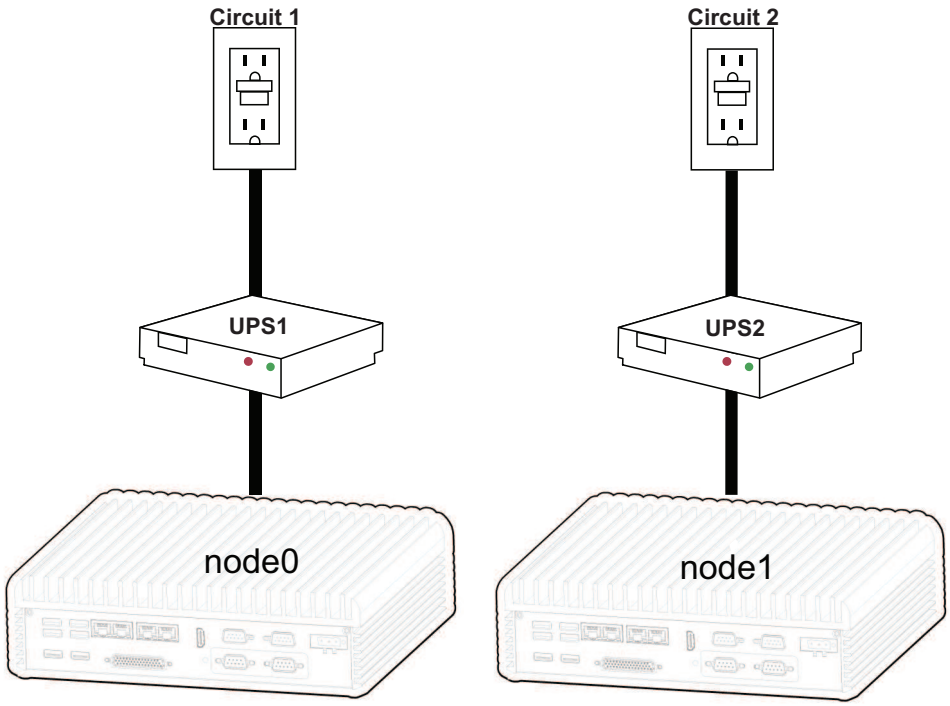

Related Topics

Power Requirements and Considerations

Deploying the System

This topic describes how to deploy a ztC Edge system. It supplements information in the Quick Deployment document for your system: for a system licensed for two nodes, see *ztC Edge 100i Systems: Quick Deployment* (R001Z); for a system licensed for one node, see the deployment document that came with the system.

To deploy a system

1. Connect the keyboard, monitor, and power to one node (Step 1 of *ztC Edge 100i Systems: Quick Deployment* (R001Z)). You can optionally connect one or two optional uninterruptible power supplies (UPS). For information, see Connecting Power.
2. The node powers on automatically. If it does not, press the power button (Step 2 of *ztC Edge 100i Systems: Quick Deployment* (R001Z)).

   When the monitor screen displays instructions, do one of the following:

- For a system licensed for two nodes, press the **[1]** key to deploy the first node.
- For a system license for one node, press the **[s]** key to deploy the single node.

3. A window appears asking you to select a keyboard map. Use the **Tab**, arrow, or **Esc** key to select one of the following:
   - **Germany - map = DE**
   - **Japan - map = JP106**
   - **USA - map = US** (the default)

   Use the **Tab** key to navigate to **OK**, and then press **Enter**.

   > **Note**: You can select or change the keyboard map after the initial deployment. For information, see Mapping Your Keyboard.

4. A message on the screen instructs you to select the method to configure the network address of this node. Use the **Tab**, arrow, or **Esc** key to select one of the following:
   - **Automatic configuration via DHCP** (the default)–Select this method to configure P1 as a dynamic IP configuration.
   - **Manual configuration (Static Address)**–Select this method to provide IP addresses for P1. A dialog box appears for you to type these values, which you obtain from your network administrator (you may have written these addresses in the **User-supplied Components** section of the deployment guide for your system):
     - IP address for this node
     - Subnet mask for this node
     - Default gateway (optional)

     If you enter invalid information, the dialog box redisplays until you enter valid information.

   Use the **Tab** key to navigate to **OK** (or **Back**), and then press **Enter**.

5. A confirmation dialog box appears. Use the arrow keys or the **Tab** key to navigate to **Save** (the default), to save the displayed values (or to navigate to **Back**, to return to the previous window). Then, press **Enter**.

   If you saved the values, a blue screen appears for several seconds.

6. For a system licensed for two nodes, a message to unpack the second node appears.

For a system licensed for one node or for two nodes, follow the messages on the screen, which also instruct you to connect network cables and, for a system licensed for two nodes, to power on the second node (Step 3 of *ztC Edge 100i Systems: Quick Deployment* (R001Z)). For additional information on the network configuration, see Connecting Ethernet Cables.

The screen continues to display various status messages for up to 15 minutes for a system licensed for two nodes, or for up to 5 minutes for a system licensed for one node.

7. The screen displays a message to connect to an IP address in a web browser (Step 4 of *ztC Edge 100i Systems: Quick Deployment* (R001Z)). Note the IP address because you will use it to log on to the ztC Console.

The monitor connected to the first node displays no more additional prompts. If you configured P1 as a dynamic IP configuration (selecting **Automatic configuration via DHCP** above for the node's network address), record its IP address as described in Recording the Management IP Address.

> **Note**: If you configured incorrect network settings (for example, you mistyped an IP address), you can correct the problem by pressing the **[1]** key to start over.

To complete the deployment, see Logging On to the ztC Console for the First Time.

Connecting Ethernet Cables

When deploying a ztC Edge system, you need to connect Ethernet cables. The following illustration shows the Ethernet cable connections for the network configuration of a system licensed for two nodes. (The keyboard and console can connect to either node0 or node1. The illustration shows the connection to node0.) On a system licensed for one node, follow the instructions (below) to connect Ethernet cables to **P1** for network0 (ibiz0) and, optionally, to **P2** for network1 (ibiz1).

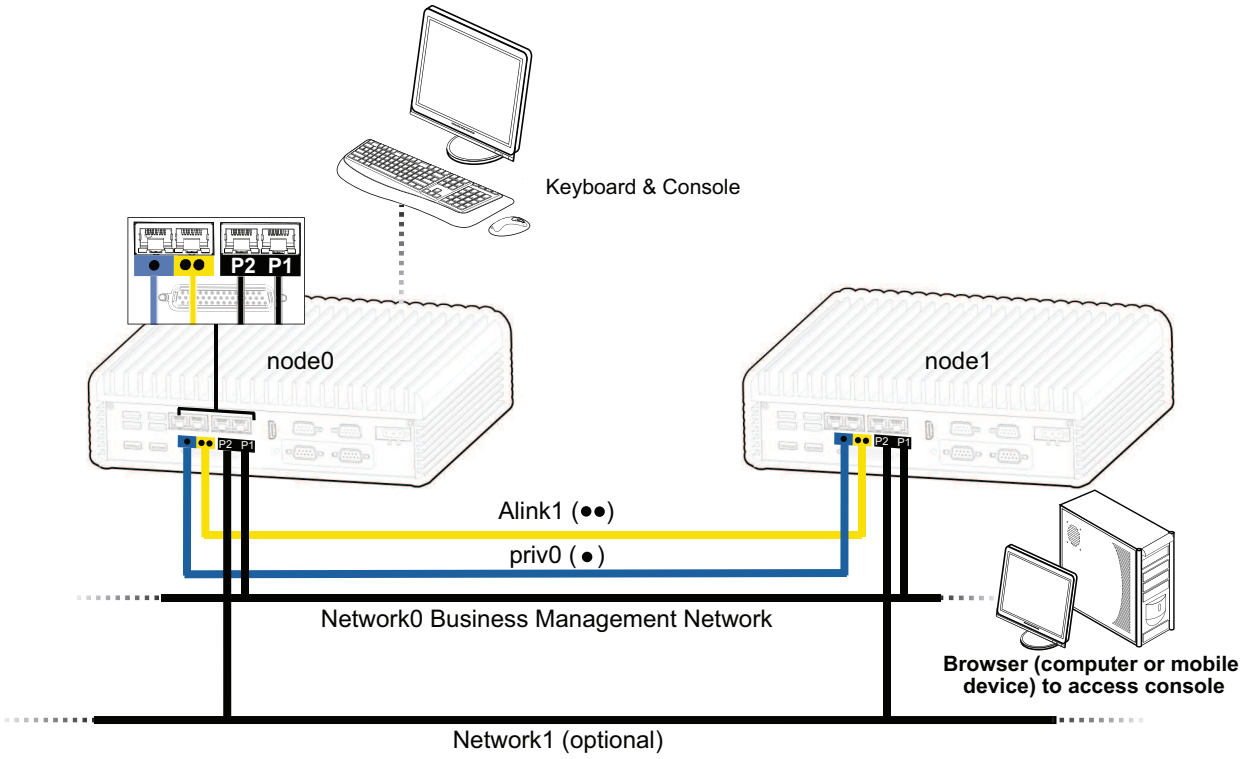

When you deploy the system (see Deploying the System), you connect:

- The blue cable for **priv0** from the blue embedded port (• or **A2** on some systems) on node0 to the same embedded port on node1.
- The yellow cable for A-Link1 from yellow embedded port (•• or **A1** on some systems) on node0 to the same embedded port on node1.

For network0 (ibiz0), you connect an Ethernet cable from **P1** on each node to a network that is accessible from the remote management computer. For the optional network1 (ibiz1), you can connect an Ethernet cable from **P2** on each node to the additional network.

Make any changes in your network (if necessary) in preparation for these connections. Then, perform the next step in Deploying the System.

Related Topics

Deployment

A-Link and Private Network Requirements

Business and Management Network Requirements

ztC Console Requirements

Mapping Your Keyboard

You can configure your keyboard for a different layout after deployment.

Supported keyboard layouts include:

Layout	Language
de	German
de-latin1	German (latin1)
de-latin1-nodeadkey	German (latin1 without dead keys)
dvorak	Dvorak
jp106	Japanese
sg	Swiss German
sg-latin1	Swiss German (latin1)
uk	United Kingdom
us	U.S. English
us-acentos	U.S. International

To configure your keyboard layout after deployment:

1. Log in to the first PM as `root`.
2. From the command line, issue the `localectl` command to configure the correct keyboard layout.

The following example configures the German keyboard layout:

```
# localectl set-keymap de
```

3. Repeat the preceding steps on the second PM, if it exists.

Related Topics

Post-Deployment Tasks

Recording the Management IP Address

Your network administrator may require the management IP address for each physical machine (PM) in order to configure the system IP address. Perform this procedure if the management network was configured to have a dynamic IP address. (Your network administrator already has this information if the management network has a static IP address.)

1. When the PM completes its installation and reboots, a screen similar to the following appears:

```
ztC Edge

     IPv4 address 10.84.52.117

     IPv6 address 3d00:feed:face:1083:225:64ff:fe8d:1b6e

     IPv6 address fe80: :225:64ff:fe8d:1b6e
```

2. Record the IPv4 address shown on the screen.
3. Give this IP address to your network administrator.

Return to Deploying the System to continue deployment.

Related Topics

Business and Management Network Requirements

Post-Deployment Tasks

After completing system deployment, you must complete several post-deployment tasks, including:

- Obtaining System IP Information
- Logging On to the ztC Console for the First Time
- Registering the System and Acquiring a Permanent License

- Configuring Required System Preferences:
  - Configuring Date and Time
  - Configuring Remote Support Settings
  - Configuring Quorum Servers
  - Specifying Owner Information
- Managing Local User Accounts
- Resolving Outstanding Alerts on the Dashboard
- Connecting a Second Business Network

In some situations, you may need to redeploy a system. For information, see Redeploying a ztC Edge System.

Obtaining System IP Information

After you deploy the system, you need the node0 IP address to log on to the ztC Console for the first time (see Logging On to the ztC Console for the First Time). To complete the initial logon procedure, you also need system IP information, which the network administrator should provide. Give the network administrator the node0 and node1 (if it exists) IP addresses (see Recording the Management IP Address), which helps the network administrator determine system IP information. The system IP address must be a static IP address. Do not use a dynamic IP address.

Related Topics

Deployment

Post-Deployment Tasks

Logging On to the ztC Console for the First Time

When deploying the system, log on to the ztC Console to accept the end-user license agreement (EULA) and to provide network information. You can also register the system and acquire a permanent license now, though you can do so later. When a system is first installed, it has a temporary license that expires within 30 days.

To log on to the ztC Console for the first time

1. From a networked PC or laptop, type the IP address of node0 (primary) into a browser address bar (Step 5 of *ztC Edge 100i Systems: Quick Deployment* (R001Z)). This IP address is the one you noted in Step 4 of *ztC Edge 100i Systems: Quick Deployment* (R001Z).

   The log-on page of the ztC Console appears.

2. Enter **admin** for the **Username** and **admin** for the **Password** (or other credentials, if provided), and then click **LOGIN**.

   The Stratus ztC Edge END USER LICENSE AGREEMENT (EULA) appears.

3. Read the EULA and then, if appropriate, click **Accept** to accept it. If you do not accept the EULA, deployment terminates.

   The **INITIAL CONFIGURATION** page appears under **Config**.

4. Under **NOTIFICATIONS**, the box for **Enable Support Notifications** is checked, by default. If you do not want the ztC Edge system to send health and status notifications to your authorized Stratus service representative, uncheck the box. You can change this setting later (see Configuring Remote Support Settings).

5. Under **SYSTEM IP**, for **Static System IP**, enter the static system IP address that you obtained from your network administrator (in *ztC Edge 100i Systems: Quick Deployment* (R001Z), see the **User-supplied Components** section). (The system IP address is sometimes referred to as the cluster IP address.)

6. Also under **SYSTEM IP**, select **DHCP** (the default) or **Static**. For **DHCP**, you do not need to provide additional information.

   If you select **Static**, the node0 static IP address that you entered during deployment appears. Provide the following values (in *ztC Edge 100i Systems: Quick Deployment* (R001Z), see the **User-supplied Components** section):

   - Primary and secondary DNS
   - NetMask
   - Gateway address for node0
   - IP address for node1 (if it exists)
   - Gateway address for node1 (if it exists)

After you have entered the network information, click **Continue**. After a short delay, the **LICENSE INFORMATION** window appears.

7. You can register the system and install a permanent license now (Step 6 of *ztC Edge 100i Systems: Quick Deployment* (R001Z)) or later. See Registering the System and Acquiring a Permanent License.
8. Click **Finish**.

The ztC Console appears and the initial logon 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-Deployment Tasks, if necessary.

Related Topics

Deployment

The ztC Console

Using the ztC Console

Registering the System and Acquiring a Permanent License

You must register a system, which includes acquiring a permanent license. When a system is first deployed, it has a temporary license that expires within 30 days. (A temporary license is displayed as **UNREGISTERED_TRIAL** for **Asset ID** in the masthead.) You can register the system when you log on to the ztC Console for the first time, or you can register it later. You can do so on a system with or without Internet access.

On a system that does not have Internet access, you need to move a file between the location of the console (which does not have Internet access) and a location with Internet access. Two methods are as follows, though other methods are possible:

- 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 Console and a location with Internet access.

Prerequisites:

- To register the system, you need the **Important Registration Information** insert from Stratus, which was shipped in the box with the system. The insert contains the ASSET ID for the ztC Edge system. If you no longer have the ASSET ID, contact your authorized Stratus service representative to obtain it.
- 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.
- **Asset ID**–Provide the ASSET ID from the Stratus registration insert.

You also need to review and accept the **Service Terms**.

Step 2: Location Information–Enter the following information:

- **Company Email** and **Retype Email**–Provide the email address of the company that is the final deployment site. Do not provide a personal email address.
- **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 PO 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 Detail–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 Stratus 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 **Stratus Customer Service Portal** at https://support.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 Stratus 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–Click **Finish** to download the license key file, which you will install on the ztC Edge system. Make note of the location where you download the file.

To register a system and acquire a permanent license

On a system with Internet access

1. If you are registering the system when logging on to the console for the first time, start with the next step. If you are registering the system after deployment, perform these steps:
    a. In the ztC Console, click **Preferences** in the left-hand navigation panel.
    b. On the **Preferences** page, click **Product License**.
2. Click the **License Check and Manual Download** bar.
3. For **Online License Registration**, click **Register** to open a new browser tab with the Stratus registration web portal. Then, complete the registration web portal steps.

    At **Step 4**, download the permanent license key file and save it to your computer.
4. In the console, click the **Install License** bar.
5. Click **Choose File** and navigate to the location where you saved the file.
6. Select the file, click **Open**, and then click **Upload** to upload the file to the system.

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

Console (which does not have Internet access) and a location with Internet access. The procedure below describes one method, though other methods are possible.

On a computer or mobile device with access to the ztC 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 Console.

2. If you are registering the system when logging on to the console for the first time, continue with the next step. If you are registering the system after deployment, perform these steps:

   a. Log on to the ztC Console.

   b. Click **Preferences** in the left-hand navigation panel.

   c. On the **Preferences** page, click **Product License**.

3. Click the **License Check and Manual Download** bar.

4. For **Offline License Registration via URL File**, 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.

5. 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 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 Console.

2. In the console, click **Preferences** in the left-hand navigation panel.

3. On the **Preferences** page, click **Product License**.

4. Click the **Install License** bar.

5. Click **Choose File** and navigate to the location where you saved the file.
6. Select the file, click **Open**, and then click **Upload** to upload the file to the system.

If you are logging on to the console for the first time, return to the final step in Logging On to the ztC Console for the First Time after you have uploaded the license.

Related Topics

Logging On to the ztC Console for the First Time

Managing the Product License

Redeploying a ztC Edge System

Redeploy a ztC Edge system if you already deployed and configured the system, but you need to reset its network settings to prepare it for deployment on a different network or subnet, possibly at a new location.

You typically redeploy a new ztC Edge system if you need to prepare it with settings and virtual machines (VMs) for an end user, but then you need to reset the network settings so the end user can deploy the system at their site for the first time (in a similar manner to using the Windows Sysprep utility to prepare a Windows system for its first end-user deployment, or Out-Of-Box Experience (OOBE)).

After configuring the system for the end user, you initiate a redeployment in the ztC Console. The system subsequently clears the system and node network settings, shuts down any VMs that are running, and powers down the system. The system retains its non-network system settings and the VMs that you configured, but it is now prepared for deployment as described in *ztC Edge 100i Systems: Quick Deployment* (R001Z), which is included with each new ztC Edge system.

To redeploy a ztC Edge system

1. Prepare the system for the end user. Configure the ztC Edge system settings and create VMs as needed. (When you redeploy the system, only the network settings will be cleared.)
2. When you are finished preparing the system, open the **Preferences** page in the ztC Console, click **IP Configuration**, and then click **Redeploy**.
3. The system clears the system and node network settings, shuts down any VMs that are running, and powers down the system.
4. The system is ready for deployment by the end user. To deploy the system, see *ztC Edge 100i Systems: Quick Deployment* (R001Z). (If needed, see Deploying the System for additional details.)

Related Topics

Deployment

The ztC Console

Using the ztC Console

Connecting a Second Business Network

When you deploy a ztC Edge system for the first time, you connect a network cable from the P1 port of each node to your existing network to create a shared business/management network called network0 (sometimes referred to as ibiz0).

If you want to add a second, dedicated business network (network1, sometimes referred to as ibiz1) after deployment, you can connect a network cable from the P2 port of each node to your existing network.

Adding a second business network may help to improve load balancing on a system with two or more VMs because you can assign the virtual machines (VMs) to separate business networks. Reducing the load on network0 can also help to improve performance because network0 carries management traffic as well as business traffic.

To connect a second business network

1. Connect a network cable from the **P2** port of each node to your existing network.
2. In the ztC Console, go to the **Networks** page.
    a. The new **network1** connection should appear within a minute or so.
    b. Verify that the new **network1** connection displays a green check.
3. Use the **Reprovision Virtual Machine** wizard to enable **network1** (and possibly disable **network0**) for each VM, as needed. For more information, see Reprovisioning Virtual Machine Resources.

Related Topics

Connecting Ethernet Cables

A-Link and Private Network Requirements

Business and Management Network Requirements

General Network Requirements and Configurations

3

Chapter 3: Using the ztC Console

The ztC Console is a browser-based interface that provides management and monitoring of an ztC Edge system from a remote management computer. For an overview of the console, see The ztC Console.

For information on pages within the ztC Console, see the following topics:

- The Dashboard Page
- The System Page
- The Preferences Page
- The Alerts History Page
- The Audit Logs Page
- The Support Logs Page
- The Physical Machines Page
- The Virtual Machines Page
- The Volumes Page
- The Networks Page
- The Virtual CDs Page
- The Upgrade Kits Page
- The Users & Groups Page

The ztC Console

The ztC Console is a browser-based interface that provides management and monitoring of an ztC Edge 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 physical machines (PMs), virtual machines (VMs), and other resources.

For information on the requirements of the remote management computer that runs the ztC Console, see ztC Console Requirements.

Using the ztC Console, you can perform a variety of administrative functions:

- Read system alerts from the Dashboard. See The Dashboard Page.
- View VM, CPU, memory, and storage statistics, and reboot or shutdown the system from the System page. See The System Page.
- Set preferences for the system, notifications (e-Alerts and SNMP configuration), and remote support (notification and access); and access administrative tools that enable you to create a secure connection. System preferences include owner information and configuration values for IP address, quorum services, date and time, etc. See The Preferences Page.
- View alerts and audit logs. See The Alerts History Page, The Audit Logs Page, and The Support Logs Page.
- Monitor, manage, and maintain resources:
    - PM status, storage (including disks), network, VMs, and USB devices: see The Physical Machines Page.
    - VM status and management tasks such as creating, importing/restoring, managing, and maintaining VMs: see The Virtual Machines Page.
    - Volumes, including their state, name, data synchronization status, size, state, and other information: see The Volumes Page.
    - Networks, including state, link condition, name, internal name, type (for example, A-Link), VMs, speed, MAC address, and network bandwidth: see The Networks Page.
    - Virtual CDs, including their state, name, size, and whether or not the VCD can be removed: see The Virtual CDs Page.
- Monitor and manage upgrade kits, and users and groups in the LIBRARY. See The Upgrade Kits Page and The Users & Groups Page.

Related Topics

Logging On to the ztC Console for the First Time

Logging On to the ztC Console

Using the ztC Console

Logging On to the ztC Console

Log on to the ztC Console to manage the ztC Edge system. Using the console, you can manage the system, including its physical machines (PMs), virtual machines (VMs), storage, and networks. You can also view alerts and logs, and perform other administrative tasks.

Note: A login session times out after one hour, if unused.

To log on to the ztC Console

1. Type the ztC Edge system's IP address or name that is a fully qualified domain name (FQDN) into a browser address bar:

   http://*IP_address*

   OR

   http://*FQDN_name*

   IP_address is the ztC Edge system's static IP address, supplied during deployment.

   FQDN_name is the FQDN corresponding to that IP address.

2. When the logon page appears, enter your **Username** and **Password**.
3. Click **LOGIN**.

Related Topics

Logging On to the ztC Console for the First Time

The ztC Console

Using the ztC Console

The Dashboard Page

The **Dashboard** page displays a summary of outstanding alerts on the ztC Edge system. To open this page, click **Dashboard** in the left-hand navigation panel.

To display additional information about outstanding alerts, click an alert symbol (for example, ) in the ztC Edge 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**, or **Ignored**, which may appear below the system diagram, depending on the alerts. The alert information includes:

- The component associated with the issue (for example, the ztC Edge system, physical machine (PM), or virtual machine (VM)).
- A description of the activity or task that requires attention.
- The reason the issue should be resolved, if available.

Resolve active alerts as soon as possible (see Resolving Outstanding Alerts on the Dashboard).

Understanding the ztC Edge System Diagram

The system diagram on the **Dashboard** page displays a graphical representation of system status. A star symbol indicates the primary PM. Alert symbols, if present, represent informational or critical alerts that require attention. Click an alert symbol to display information about the alert.

Related Topics

The Physical Machines Page

The System Page

The Virtual Machines Page

Resolving Outstanding Alerts on the Dashboard

After completing system deployment, resolve any outstanding alerts that appear on the Dashboard page.

To resolve outstanding alerts

On the ztC Console Dashboard page, view any alerts listed in the lower portion of the page. 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**, then enable support notification service.
- Click **Ignore** (beneath the **Action** column) to ignore the alert and remove it from the list. Minor alerts can be ignored rather than resolved. Clicking **Ignore** hides the alert.

To restore the ignored alert to the list, click **Ignored**, above the alerts list, and then **Restore**, under the **Action** column.

Related Topics

The Dashboard Page

The System Page

The **System** page displays information about the ztC Edge system, and enables you to reboot or shut down the system. The page also displays statistics and resource allocations for the ztC Edge system. To open this page, click **System** in the left-hand navigation panel.

You can use the **System** page for administrative tasks including:

- Rebooting the System
- Shutting Down the System

To power on the system (at the physical console of the PMs), see Powering On the System.

You perform many other administrative tasks on the ztC Edge system using the ztC Console. For information, see The ztC Console.

Viewing statistics

The **System** page contains these sections, which display information and statistics of system usage as well as of PMs and VMs:

- ***system name***–Circle graphs indicate the system's CPU allocation, memory allocation, disk (R/W), and network utilization.
- **Node0** and **Node1** (if it exists)–Circle graphs indicate each node's CPU utilization, memory utilization, disk utilization, and network utilization. For disk utilization and network utilization, you can select the logical disk or the network whose statistics you want to display.

Related Topics

Using the ztC Console

Powering On the System

Power on the ztC Edge system at the physical console of each physical machine (PM), or node. Doing so performs an orderly startup by first booting the system software and then starting the virtual machines (VMs) on the system. (To power off a system, see Shutting Down the System.)

Caution: If you are powering on the system for the first time to deploy it, follow the instructions in *ztC Edge 100i Systems: Quick Deployment* (R001Z). (If needed, see Deploying the System for additional details.)

Note: If a PM loses power because you disconnect the power cord or AC mains power is lost, each PM in a dual-node system, and the one PM in a single-node system, is set to power on automatically as soon as power is restored. The system software and VMs restart automatically.

To power on a ztC Edge system

1. Ensure that all the required network cables are connected. In a dual-node system, check that the network cables are connected to both PMs.
2. Press the power button on the front panel of the PM(s) in the system.
3. Ensure that the **PWR** LED on the front panel of the PM(s) is lit.

Related Topics

The ztC Console

The System Page

Using the ztC Console

Rebooting the System

Reboot the ztC Edge system using the ztC Console to safely restart both PMs without affecting VMs. On a single-node system, reboot the system only during a planned maintenance period, since the Virtual Machines are shut down and restarted during the reboot.

Caution: Rebooting the ztC Edge system by any method other than following (for example, rebooting from the PMs individually) may result in data loss.

Note: You can reboot a system licensed for two PMs only if both PMs are running, healthy, and not in maintenance mode. You can reboot a system licensed for one PM only if the PM is running, healthy, and not in maintenance mode.

Prerequisite: On a system licensed for two PMs, confirm that both PMs are running before rebooting. On a single-node system, confirm that the one PM is running before rebooting.

To reboot the ztC Edge system

1. Select **System** in the left-hand navigation panel.
2. Click the **Reboot** button. A message appears, asking you to confirm the reboot. Click **Yes** to continue.

   Rebooting can take up to 15 minutes. You can observe the process in the **Dashboard** and the masthead of the ztC Console. The system's PMs sequentially enter and then exit maintenance mode (for information on maintenance mode, see Maintenance Mode).
3. Verify that the PMs restart and that all VMs continue running as expected.

After you initiate a reboot, a message in the masthead shows the status of the reboot. If necessary, you can cancel the reboot by clicking **Cancel Reboot** in the masthead.

Caution: If you cancel a reboot, the system is left in its current state and you need to manually restore it to a healthy state.

Related Topics

The ztC Console

The System Page

Using the ztC Console

Shutting Down the System

Use the ztC Console to shut down the ztC Edge system. Doing so performs an orderly shutdown by first shutting down the virtual machines (VMs) and then the physical machines (PMs). Use only this method to shutdown the ztC Edge system. Before shutting down, make sure both PMs of a system licensed for two nodes are running, and the one PM of a system licensed for one node is running.

Cautions:

1. Shutting down the ztC Edge system takes the VMs offline, so shutdown the system only during a planned maintenance period.
2. Shutting down the ztC Edge system by any other method (for example, removing power from both PMs individually) may result in data loss.

Note: When you shut down the system, standby power remains on for lights-out management unless you disconnect the power cord or the AC mains power is switched off.

To shut down the ztC Edge system

1. On systems licensed for two nodes, confirm that both PMs are running so that the disks can synchronize between nodes.
2. Select **System** in the left-hand navigation panel.
3. Click the **Shutdown** button. A warning appears: *It will shut down the entire system and stop one or more VMs!* Click **Yes** to shutdown or **No** to cancel the shutdown. After clicking **Yes**, a second warning appears, asking you to confirm the shutdown. Click **Yes** (again) to shutdown or **No** to cancel the shutdown.

You can observe some of the shutdown process in the **Dashboard** and the masthead of the ztC Console as the system's PMs sequentially enter maintenance mode (for information on maintenance mode, see Maintenance Mode). When the system shuts down completely, though, the ztC Console is unavailable and the masthead displays **Lost Communication**.

After the system shuts down, you lose the connection to the console. If the ztC Edge system cannot shutdown completely, a VM may not be shutting down properly. Do one of the following to shut down the VM:

- Use the VM console or a remote desktop application to log on to the VM. Use operating system commands to shut down the VM.
- Log on to the ztC Console. Click **Virtual Machines** in the left-hand navigation panel, select the VM, and then click **Power Off**.

Related Topics

Managing the Operation of a Virtual Machine

The ztC Console

The System Page

Using the ztC Console

The Preferences Page

The **Preferences** page enables you to configure ztC Edge system settings. To open this page, click **Preferences** in the left-hand navigation panel.

The following table lists and describes the preferences.

Preference	Description
System	
Owner Information	Allows you to specify and then view the name and contact information for an ztC Edge system administrator. This information is also provided in response to Simple Network Management Protocol (SNMP) requests. See Specifying Owner Information.
Product License	Allows you to view and manage the ztC Edge product license. See Managing the Product License.
IP Configuration	Allows you to view and specify the Internet Protocol (IP) address and network settings for the system; and to redeploy a system. See Configuring IP Settings.
Quorum Servers	Allows you to view existing and new Quorum servers. Quorum servers provide data integrity assurances and automatic restart capabilities for specific failures in the ztC Edge environment. See Quorum Servers and Configuring Quorum Servers.
Date & Time	Allows you to view the system time, specify values for Network Time Protocol (NTP) (recommended), or to manually set the time and date on the sys-

Preference	Description
	tem. See Configuring Date and Time.
Administrative Tools	
Secure Connection	Allows you to enable only HTTPS connections to the system. See Configuring Secure Connections.
VM Device Configuration	Allows you to disable or enable insertion of virtual CDs (VCDs) in all VMs or attachment of USB devices to all VMs. See Configuring VM Devices.
Notification	
e-Alerts	Allows you to enable email alerts (e-Alerts) for system administrators. See Configuring e-Alerts.
SNMP Configuration	Allows you to enable Simple Network Management Protocol (SNMP) requests and traps for remote system monitoring. See Configuring SNMP Settings.
Remote Support	
Support Configuration	Allows you to configure remote access and notifications. Remote access enables your authorized Stratus service representative to log on to the system remotely for troubleshooting. When enabled, the ztC Edge system can send notifications to your authorized Stratus service representative about problems with the system. See Configuring Remote Support Settings.
Proxy Configuration	Allows you to configure proxy settings for the ztC Edge system if your organization requires a proxy server to access the Internet and you have a service agreement with Stratus or another authorized ztC Edge service representative. The Stratus Redundant Linux software uses proxy server information for support notification messaging and remote support access features. See Configuring Internet Proxy Settings.

Related Topics

The ztC Console

Using the ztC Console

Specifying Owner Information

Specify the name and contact information for an administrator or owner of the ztC Edge system to make this information available for support purposes.

This contact information is available in the ztC Console and provided in response to Simple Network Management Protocol (SNMP) requests.

To specify system owner information

1. Click **Preferences** in the left-hand pane.
2. On the **Preferences** page, click **Owner Information**.
3. Supply information in the **Full Name**, **Phone Number**, **Email**, and **Site Address** fields.
4. Click **Save**.

Related Topics

The Preferences Page

The ztC Console

Managing the Product License

Manage the product license for the system by:

- Acquiring a permanent license during or after deployment.
- Checking the status of an existing license and, if necessary downloading a new license.
- Viewing current license information such as status and expiration date.

When a system is first installed, it has a temporary license that expires within 30 days. (A temporary license is displayed as **UNREGISTERED_TRIAL** for **Asset ID** in the masthead.) You must register the system, which includes acquiring a permanent license. You can register the system immediately after the initial deployment or later. For information on registering the system, see Registering the System and Acquiring a Permanent License.

Once a system has a permanent license, it checks with the license server for updates every 24 hours, if the system has an Internet connection. If a system does not have Internet access, you can still update the license and check its status. To do so, you need to move a file between the location of the ztC Console (which does not have Internet access) and a location with Internet access. Two methods are as follows, though other methods are possible:

- 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 Console and a location with Internet access.

Choose the menu below (click drop-down, if applicable) for the procedure that is appropriate for your needs.

To check the status of a license

If the system has Internet access, use the following procedure If the system does not have Internet access, use the ***On a system without Internet access*** procedure. If you need to update a license manually, see ***To update a new license manually***.

1. In the ztC Console, click ***asset_ID*** (of **Asset ID: *asset_ID***) in the masthead.

   Alternatively, click **Preferences** in the left-hand navigation panel of the console, and then:

   a. On the **Preferences** page, click **Product License**.

   b. Click the **License Check and Manual Download** bar.

   c. For **Online License Check**, click **Check License Now**.

2. The console displays the status of the license (date format varies, based on location):

STATUS	License is activated and does not expire.
LAST CHECK	*day, month dd*, 20*yy*, *time*
SERVICE EXPIRATION	*day, month dd*, 20*yy*, *time*
ASSET ID	*asset_ID*
PRODUCT UUID	*xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx*
ALSR Allowed	*Yes_or_No*
Guest Monitoring Allowed	*Yes_or_No*

To update a new license manually for a registered system

On a registered system with an Internet connection, the license is updated automatically. You can also, if necessary, update a license manually.

On a system with Internet access

1. In the console, click **Preferences** in the left-hand navigation panel.
2. On the **Preferences** page, click **Product License**.
3. Click the **License Check and Manual Download** bar.
4. For **Online License Check**, click **Check License Now**.

   The license is verified automatically. If the system needs a new license, the new license is automatically applied, with an updated service end date and any new features purchased.

On a system without Internet access

Use the procedure below to check a license and, if necessary, acquire a new license manually on a registered system that does not have Internet access. You need to move a file between the location of the ztC Console (which does not have Internet access) and a location with Internet access. The procedure below describes one method, though other methods are possible.

On a computer or mobile device with access to the ztC 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 Console.
2. Log on to the ztC Console.
3. Click **Preferences** in the left-hand navigation panel.
4. On the **Preferences** page, click **Product License**.
5. Click the **License Check and Manual Download** bar.
6. For **Offline License Download via URL File**, click **Download URL File** and save the file to the USB flash drive or mobile device. If using a USB flash drive, remove it. 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.
3. A web browser opens and the Stratus license server checks the status of the license file. If necessary, a new license .key file is automatically downloaded. If using a USB flash drive, copy the new license .key file to it, and then remove the USB flash drive.
4. Return to the location with access to the console.

On a computer or mobile device with access to the ztC 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 Console.
2. In the console, click **Preferences** in the left-hand navigation panel.
3. On the **Preferences** page, click **Product License**.
4. Click the **Install New License** bar.
5. Click **Choose File**, and navigate to the location where you saved the file.
6. Select the file, click **Open**, and then click **Upload** to upload the file to the system.

Related Topics

The ztC Console

The Preferences Page

Using the ztC Console

Configuring IP Settings

Configure Internet Protocol (IP) settings for the ztC Edge system to set or modify the IP address of the system and nodes as well as values for applicable settings such as network mask, gateway address, and Domain Name System (DNS) server. (You also modify network settings when redeploying a system using the **Redeploy** button, as described in Redeploying a ztC Edge System.)

During deployment and post-deployment, you configure IP addresses for the system. For a system licensed for two nodes, you configure three IP addresses: one for the system and one for each node (node0 and node1). For a system licensed for one node, you configure two IP addresses: one for the system and one for

the node (node0). You can change the IP addresses and other IP settings after deployment using the appropriate procedure below. You must specify a static IPv4 address for the ztC Edge system.

Warnings:

1. Do not change the IP configuration settings, especially on systems with running VMs, without the advice and knowledge of your network administrator. Doing so could make the system and all its VMs inaccessible.
2. If you change the **Static System IP** address, any MAC addresses automatically assigned to the VMs will change when the VMs reboot, because the Stratus Redundant Linux software generates MAC addresses for the VMs based on the system IP address. To prevent changes to the MAC address for a VM (for example, to support software applications that are licensed on a MAC-address basis), set a persistent MAC address as described in Assigning a Specific MAC Address to a Virtual Machine.
3. You must use the ztC Console to change IP addresses. Do not use Linux tools.

Notes:

1. The procedure you use to configure IP settings depends on whether the ztC Edge system stays on the same subnet or moves to a new subnet. If you need to move a ztC Edge system to a new subnet, *redeploy* the system to clear its network settings before moving it, as described in Redeploying a ztC Edge System.
2. Changing IP settings for a new subnet typically includes changing the node's physical network connections (for example, disconnecting and then re-attaching network cables if moving the PMs). Before you disconnect cables from nodes, you must shut down the nodes. For this, you have the option of using the **Save and Shutdown** button in the **IP Configuration** section of the **Preferences** page.
3. In a system licensed for one node, the **IP Configuration** page displays settings for only one node.

To change the system and/or node IP settings with the system on same subnet

The ztC Edge system and all virtual machines (VMs) continue to run throughout this procedure; however, the ztC Console briefly loses its connection to the system if you change the system IP address. You can access

the ztC Console at the new system IP address within 1-2 minutes. (You can change node IP addresses on each node, individually, but the console connection is not lost.)

1. Click **Preferences** in the left-hand navigation panel, to open the **Preference** page.
2. Click **IP Configuration**.
3. In the **Static System IP** box, type the static system IP address that you obtained from your network administrator.
4. Click the **Static** button and type valid, unique values for **Primary DNS** and **Secondary DNS**.
5. Verify that the displayed **NetMask** value is correct.
6. For **Node0** and **Node1** (if it exists), enter appropriate values for **IP Address** and **Gateway IP**.
7. Click **Save** to save the values (or click **Reset** to restore previous values).

   If you have changed the system IP address, the **System IP has been updated** message box appears. After a brief delay, the browser redirects automatically to the new system IP address.

Related Topics

Deployment

Obtaining System IP Information

Logging On to the ztC Console for the First Time

The Preferences Page

The ztC Console

Using the ztC Console

Configuring Quorum Servers

When you log on to the ztC Edge system for the first time, configure quorum servers.

Prerequisites:

1. You must have a system licensed for two nodes in order to configure a quorum server.
2. Before you configure quorum servers, read Quorum Servers and Creating an ALSR Configuration (which discusses quorum servers).

Notes:

1. For a VM to recognize quorum server configuration changes, you must reboot the VM by shutting it down and then restarting it. See Shutting Down a Virtual Machine and Starting a Virtual Machine.
2. Windows Updates on a quorum server can interrupt the server's operation, which affects fault-recovery behavior. On quorum servers, you should schedule Windows Updates during a maintenance period or disable Windows Updates.

To configure quorum servers

1. Click **Preferences** in the left-hand navigation panel, to open the **Preferences** page.
2. Click **Quorum Servers**.
3. Click **Add Quorum Server**.
4. In the **Add Preferred Quorum Server** dialog box, enter the following values (if a preferred quorum server already exists, the **Add Alternate Quorum Server** dialog box appears):
   - **DNS or IP Address**–Type the fully-qualified **DNS** host name or **IP address** for the preferred quorum server.
   - **Port** (the default value is 4557)–Type the port number if it is different from the default.

   Click **Save** to save the values.
5. Repeat steps 4 and 5 to configure a second, alternate quorum server. Stratus recommends configuring two quorum servers.
6. To enable quorum service, select the **Enabled** check box and click **Save**.

To remove a quorum server

Caution: If you remove the preferred quorum server, the alternate quorum server becomes the preferred quorum server. If no alternate quorum server exists, removing the preferred quorum server automatically disables quorum service.

1. Navigate to the **Preferences** page of the ztC Console.
2. Click **Quorum Servers**.

3. Locate the entry for the quorum server you want to remove.
4. In the right-most column, click **Remove**.

> **Note**: If a VM is using the quorum server that you are removing, you must reboot the VM so that it no longer recognizes the quorum server, which allows the removal process to finish.

Related Topics

The ztC Console

The Preferences Page

Using the ztC Console

Configuring Date and Time

When you log on to the ztC Edge system for the first time, configure the date and time to enable the Network Time Protocol (NTP) service. Using the NTP service automatically sets the system clock and ensures that it does not drift from the actual time.

> **Caution**: When you change the date and time settings, the primary physical machine (PMs) may reboot and the secondary PM (if it exists) may shutdown if system time has drifted from actual time. All virtual machines (VMs) are stopped and business processing is interrupted until the reboot is complete.

> **Note**: The clock swaps between time zones whenever VMs migrate or restart. To ensure that the time zone in VMs does not change:
>
> - Set the time zone in all VMs to correspond to the time zone configured for the ztC Edge system.
> - Configure all VMs to use the same NTP servers as those configured for the ztC Edge system.

To configure date and time settings

1. Click **Preferences** in the left-hand navigation panel, to open the **Preferences** page.
2. On the **Preferences** page, click **Date & Time**.

3. In the **Date & Time** display, the default setting for **Configure Time Zone** is **America, New York**. Select a time zone appropriate for your location, if necessary.
4. Select one of the following for **Configure Date and Time**:
   - **Automatically (recommended)** enables NTP service. Type NTP server addresses in the text area, one per line. Specifying multiple NTP servers provides redundancy.
   - **Manually** allows you to manually enter settings.

   > **Note**: If you configure time manually, the ztC Edge system's time may drift from actual time.
5. Click **Save** (or click **Reset** to restore the previously-saved values).

If the system requires a reboot because of time drift, a message appears in the ztC Console masthead telling you that the system will reboot. In this case, the primary physical machine (PM) reboots and the secondary PM (if it exists) shuts down. While the primary PM reboots, you lose your connection to the ztC Console. When the reboot is complete, the PM re-establishes a connection to the console and you receive an alert telling you to restart the secondary PM.

Related Topics

The ztC Console

The Preferences Page

Using the ztC Console

Configuring Secure Connections

For security, you may want the ztC Edge system to allow only HTTPS connections. By default, the system allows HTTPS and HTTP connections.

> **Note**:
> When you activate or deactivate the check box next to **Enable HTTPS Only / Disable HTTP** in the procedure below and click **Save**, the system automatically logs you out of the ztC Console and you must log in again,

When HTTPS connections are enabled, you can use a script to install a custom certificate on the host machine. See To install a custom certificate.

To enable only HTTPS connections

1. Click **Preferences** in the left-hand navigation panel, to open the **Preferences** page.
2. Under **Administrative Tools**, click **Secure Connection**.
3. Activate the check box next to **Enable HTTPS Only / Disable HTTP**.
4. Click **Save**.

   The system automatically logs you out of the ztC Console, redirects the browser to the HTTPS login page, and you must log in again.

If the system allows only HTTPS connections and you want to allow HTTP connections, you need to deactivate the check box.

To enable HTTP and HTTPS connections

1. Click **Preferences** in the left-hand navigation panel, to open the **Preferences** page.
2. Under **Administrative Tools**, click **Secure Connection**.
3. Deactivate the check box next to **Enable HTTPS Only / Disable HTTP**.
4. Click **Save**.

   The system automatically logs you out of the ztC Console and redirects the browser to the HTTPS login page. To access the HTTP login page, you manually replace **https** with **http** in the browser's address bar, and then you can log in.

To install a custom certificate

To install a custom certificate, use the `certificate_installer` script. Using this script, you can install a custom SSL certificate, recover a previously used or build-in certificate, and display information about a certificate currently in use or previously used, as follows:

- Install a custom certificate (non HTTPS-only mode):
  i. Copy a certificate to the `/tmp` folder of the host machine.
  ii. Issue the following command:

  ```
  certificate_installer install -c /tmp/server.crt -k
  /tmp/server.key
  ```

- Install a custom certificate (HTTPS-only mode):

i. Copy a certificate to the `/tmp` folder of the host machine.

ii. Issue the following command:

```
certificate_installer install -c /tmp/server.crt -k
/tmp/server.key -f
```

- Recover the custom certificate to the previously used one:

```
certificate_installer recover -p
```

- Recover the custom certificate to the built-in one:

```
certificate_installer recover -b
```

- List information about the currently used certificate:

```
certificate_installer list -c
```

- List information about the previously used certificate:

```
certificate_installer list -p
```

The `certificate_installer` script

Usage

```
certificate_installer [command command_options] [script_options]
```

Commands and Command Options

`install` *command_options*	Installs the custom certificate. Command options are: • `-c, --cert=`*`certificate_path`*: The path where the certificate is saved. • `-k, --key=`*`private_key_path`*: The path where the key is saved. • `-f, --[no-]force`: Force replacing the SSL certificate in use.
`recover` *command_options*	Recovers the custom certificate. Command options are: • `-b, --[no-]built-in` (the default): Recover to the built-in certificate. • `-p, --[no-]previous`: Recover to the previously used certificate
`list` *command_options*	Lists the custom certificate(s). Command options are: • `-a, --[no-]all` (default): List all SSL certificates on host machine. • `-c, --[no-]current`: List the currently used certificate. • `-p, --[no-]previous`: List the previously used certificate. • `-L, --location=`*`location`*: Show information of a certificate at a specified location.

Script Options

`-v, --[no_]verbose`	In verbose mode, the script displays all information.
`-l, --log=`*log_file*	Prints logs to the file *log_file* instead of to `STDOUT`.

Related Topics

The ztC Console

The Preferences Page

Using the ztC Console

Configuring VM Devices

Configure VM devices to disable or enable insertion of virtual CDs (VCDs) in all VMs or attachment of USB devices to all VMs. By default, these VM devices can be inserted and attached. Use **VM Device Configuration** on the **Preferences** page to change the configuration.

When VM devices are enabled (the default) for insertion or attachment, you can insert VCDs in all VMs or attach a USB device to VMs. When VM devices are disabled for insertion or attachment, you cannot insert or attach these devices.

To disable insertion or attachment of VM devices

1. Click **Preferences** in the left-hand navigation panel, to open the **Preferences** page.
2. On the **Preferences** page, click **VM Device Configuration** beneath **Administrative Tools**.
3. Activate the check box for one or both of the following:
   - **Disable insertion of CDs on all VMs**–Activate the check box to disable inserting CDs in VMs.
   - **Disable attachment of USB devices to all VMs**–Activate the check box to disable attaching USB devices to VMs.
4. Click **Save**.

To enable insertion or attachment of VM devices

1. Click **Preferences** in the left-hand navigation panel, to open the **Preferences** page.
2. On the **Preferences** page, click **VM Device Configuration** beneath **Administrative Tools**.
3. Deactivate the check box for one or both of the following:
   - **Disable insertion of CDs on all VMs**–Deactivate the check box to enable inserting CDs in VMs.
   - **Disable attachment of USB devices to all VMs**–Deactivate the check box to enable attaching USB devices to VMs.
4. Click **Save**.

Related Topics

Inserting a Virtual CD

Attaching a USB Device to a Virtual Machine

The Preferences Page

Using the ztC Console

Configuring e-Alerts

Configure email alerts (e-Alerts) to enable the ztC Edge system to send email to system administrators whenever the system detects an event requiring administrator attention.

To enable e-Alerts

1. Click **Preferences** in the left-hand navigation panel, to open the **Preference** page.
2. Under **Notification**, click **e-Alerts**.
3. Click the **Enable e-Alerts** 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** (optional)–Enter the port number to use when sending e-Alerts. If no port number is specified, the default SMTP port 25 will be used. (For additional information on all ports, including the SMTP port, access the Knowledge Base to search for the article *TCP and UDP ports used by ztC Edge* (KB-9357). See Accessing Knowledge Base Articles.)

- **e-Alerts Language**–Select a language from the pull-down menu.
- **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 Edge 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 pull-down menu for the encryption protocol that the SMTP server requires:
  - **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.
  - **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.
- **Enable Authentication**–Click this box if the SMTP server requires authentication to send email. Then, type the **Username** and **Password** for the SMTP account.
- **List of Recipients** (required)–Enter email addresses for all e-Alert recipients.

4. Click **Save** (or click **Reset** to restore the previously-saved values).

> **Note**: When you enable or update the e-Alert configuration, generate a test alert to confirm that you receive the alerts.

To generate a test alert

Click **Generate Test Alert**. The Stratus Redundant Linux 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); and Support Configuration sends a notification to your authorized Stratus service representative, if configured (see Configuring Remote Support Settings). Watch the Alerts History log (see The Alerts History Page) for delivery status.

You can also test e-Alerts by putting the secondary physical machine into maintenance mode (see Maintenance Mode), and then removing it from maintenance mode. Verify that you receive e-Alerts for both maintenance mode events.

Related Topics

The ztC Console

The Preferences Page

Using the ztC Console

Configuring SNMP Settings

Configure Simple Network Management Protocol (SNMP) settings for the ztC Edge system to allow SNMP management applications to remotely monitor your systems. (SNMP information pertains only to systems and not individual PMs.) 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 Stratus Redundant Linux software. These MIBs include a system-specific MIB that is a collection of objects describing the ztC Edge system. You can download a copy of the MIB file from the **Drivers and Tools** section of the **Downloads** page at https://www.stratus.com/services-support/downloads/?tab=ztcedge.
- SNMP trap–A message initiated by one of the nodes in the ztC Edge system after an event such as an alert that is then sent to an identified list of recipients, typically a network management station (NMS).

Follow the appropriate procedure to enable SNMP requests or traps.

To enable SNMP requests

You can allow no users (**Restricted**) or any user using the default `public` community (**Unrestricted**) to send requests. Or you can customize the list of users that can send requests by editing the `snmpd.conf` file.

To customize the list

1. Log in to the host.
2. Manually edit the standard `/etc/snmp/snmpd.conf` file on both nodes.
3. Save the file.
4. Restart the `snmpd` process on each node by entering the command **systemctl restart snmpd**.

To enable SNMP requests

1. Click **Preferences** in the left-hand navigation panel, to open the **Preference** page.
2. Under **Notification**, click **SNMP Configuration**.
3. Activate the check box next to **Enable SNMP Requests**.

   The selections in Step 4 appear as well as a read-only display of the `snmpd.conf` file.
4. Select one of the following:

   Restricted (the default)–Allows no users to send requests.

   Unrestricted–Allows public access.

   Customized (available when `snmpd.conf` has been edited)–Allows customized access.
5. Click **Save**. (Or click **Reset** to restore the previously-saved values.)

To enable SNMP traps

> **Note**: When you enable or modify the SNMP trap settings, generate a test alert to confirm that traps are received.

1. Click **Preferences** in the left-hand navigation panel, to open the **Preference** page.
2. Under **Notification**, click **SNMP Configuration**.
3. Activate the check box next to **Enable SNMP Traps**.
4. Type the name of the SNMP **Community**, or keep the default (**public**).
5. Next to **List of Recipients for SNMP traps (Version 2c)**, type the IP address or host name for each SNMP version 2 recipient, one per line.
6. Next to **List of Recipients for SNMP traps (Version 1)**, type the IP address or host name for each SNMP version 1 recipient, one per line.
7. Click **Save**. (Or click **Reset** to restore the previously saved values.)
8. Configure your organization's firewall to allow SNMP operations, which enables SNMP management systems to receive alerts from and send traps to the ztC Edge system. To do so, configure your organization's firewall to open the SNMP port:

   Message Type: SNMP

   Protocol: SNMP

   Port: 161 (Get/Walk) 162 (Traps)

9. Generate a test alert by clicking **Generate Test Alert**.

   The Stratus Redundant Linux 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); and Support Configuration sends a notification to your authorized Stratus service representative, if configured (see Configuring Remote Support Settings). Watch the Alerts History log (see The Alerts History Page) for delivery status.

Related Topics

SNMP

The ztC Console

The Preferences Page

Using the ztC Console

Configuring Remote Support Settings

When you log on to the ztC Edge system for the first time, configure support configuration settings that enable the ztC Edge system to send support notifications (alerts) to your authorized Stratus service representative when an event requires attention.

To configure support configuration settings

Note: When you enable or modify settings for **Enable Remote Support Access** or **Enable Notifications**, generate a test alert to confirm that your authorized Stratus service representative can receive system health messages from your system.

1. Click **Preferences** in the left-hand navigation panel, to open the **Preferences** page.
2. Under **Remote Support**, click **Support Configuration**.
3. Modify the settings, as appropriate for your system:
   - **Enable Remote Support Access** allows your authorized Stratus service representative to remotely connect to the ztC Edge system for troubleshooting purposes. Note that you can enable and then disable this setting, as needed.
   - **Enable Notifications** allows the ztC Edge system to send health and status notifications to your authorized Stratus service representative.

- **Enable Support Notifications** sends an alert for any event that requires attention. It also sends a periodic "heartbeat" call-home message to your authorized Stratus service representative.
- **Enable Periodic Reporting** sends a daily summary of system information to help improve product and service quality.

4. Click **Save** (or click **Reset** to restore the previously saved values).
5. Configure your organization's firewall to allow support messages.

To configure your firewall to allow support messages

Use the following information to configure your organization's firewall to allow communication with your authorized Stratus service representative:

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: Dial-In
Protocol: TCP
Port: 443, Default proxy port: 3128 (You can change the default proxy port number.)
Stratus support server address: *.ecacsupport.com

Message Type: e-Alert
Protocol: SMTP
Port: 25

(For additional information on TCP and UDP ports, access the Knowledge Base to search for the article *TCP and UDP ports used by ztC Edge* (KB-9357). See Accessing Knowledge Base Articles.)

To enable SNMP management systems to receive alerts and send traps to the ztC Edge system, configure the firewall for the following:

Message Type: SNMP

Protocol: SNMP

Port: 161 (Get/Walk) 162 (Traps)

6. Generate a test alert.

   To generate a test alert

   Click **Generate Test Alert**. The Stratus Redundant Linux software generates a test alert and Support Configuration sends a notification to your authorized Stratus service representative; e-Alerts send a sample email with the subject "Test Alert" to all email recipients of e-Alerts, if configured (see Configuring e-Alerts); and SNMP sends traps to recipients of SNMP traps, if configured (see Configuring SNMP Settings). Watch the Alerts History log (see The Alerts History Page) for delivery status. A subsequent alert will be generated if the support notification fails.

Related Topics

The ztC Console

The Preferences Page

Using the ztC Console

Configuring Internet Proxy Settings

Configure proxy settings for the ztC Edge system if your organization requires a proxy server to access the Internet and you have a service agreement with Stratus or another authorized ztC Edge service representative.

A proxy server provides a secure bridge between the ztC Edge system and the Internet. Stratus Redundant Linux 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. Click **Preferences** in the left-hand navigation panel, to open the **Preferences** page.
2. Under **Remote Support**, click **Proxy Configuration**.
3. To enable proxy service, click the **Enable Proxy** box.
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**.

7. Click **Save** (or click **Reset** to restore the previously-saved values).

Related Topics

The ztC Console

The Preferences Page

Using the ztC Console

The Alerts History Page

The **Alerts History** page displays messages about events on the ztC Edge system.

To open the **Alerts History** page, click **Alert History** in the left-hand navigation panel of the ztC Console. (To view a log of user activity on the ztC Edge system, see The Audit Logs Page.)

Note: Support notification alerts, e-Alerts, and SNMP traps are generated only when you enable them in the ztC Console console. For information see:

- Configuring Remote Support Settings
- Configuring e-Alerts
- Configuring SNMP Settings

To view alert information, scroll through the alerts, which are, by default, listed in reverse chronological order. Click an alert to display the time the alert occurred as well as information about the problem and resolution (if available), and whether **Support Notifications**, an **e-Alert**, or an **SNMP Trap** was sent for this alert. (You can also display alert information using `snmptable`; see Obtaining System Information with snmptable.)

To remove an alert, select it and click **Remove**.

To remove all of the alerts, click **Purge All**.

Related Topics

The ztC Console

Using the ztC Console

The Audit Logs Page

The **Audit Logs** page displays a log of user activity in the ztC Console. To open this page, click **Audit Logs** in the left-hand navigation panel. (To display information about events on the ztC Edge system, see The Alerts

History Page.)

To view log information, scroll through the log entries, which are, by default, listed in reverse chronological order. The information includes:

- **Time**–The date and time of the action.
- **Username**–The name of the user that initiated the action.
- **Originating Host**–The IP address of the host on which the ztC Console was running.
- **Action**–The action performed in the ztC Console.

You can also display information about audit logs using `snmptable` (see Obtaining System Information with snmptable).

Related Topics

The ztC Console

Using the ztC Console

The Support Logs Page

The **Support Logs** page enables you to generate diagnostic files, which include the ztC Edge system's log files and configuration information at a particular moment in time. This information enables your authorized Stratus service representative to resolve an issue with the system.

For additional information, see:

- Creating a Diagnostic File
- Deleting a Diagnostic File
- Uploading a Diagnostic File to Customer Support

Related Topics

The ztC Console

Using the ztC Console

The Preferences Page

Creating a Diagnostic File

Diagnostic files provide the ztC Edge system's log files and configuration information at a particular moment in time. You create a diagnostic file to help your authorized Stratus service representative resolve issues with

the system.

> **Note**: Stratus Redundant Linux software allocates a fixed amount of storage space for diagnostic files. If sufficient space is not available when you create a diagnostic file, the system will delete previously created files.

To create diagnostic files

1. Click **Support Logs** in the left-hand navigation panel, to open the **Support Logs** page.
2. Click **Generate Diagnostic File**.
3. Upload the file to your authorized Stratus service representative, as described in Uploading a Diagnostic File to Customer Support.

Related Topics

The ztC Console

The Preferences Page

Using the ztC Console

Uploading a Diagnostic File to Customer Support

Upload a diagnostic file to the Stratus ztC Edge Customer Support web site to help resolve an issue with the system. (To create a diagnostic file, see Creating a Diagnostic File.)

To upload a diagnostic file to Customer Support

1. Click **Support Logs** in the left-hand navigation panel, to open the **Support Logs** page.
2. Do one of the following:
   - If the ztC Edge system has Internet connectivity, upload the diagnostic file directly to the Stratus ztC Edge Customer Support web site by clicking **Upload**. If the upload succeeds, a message appears, confirming that the diagnostic file was uploaded successfully.
   - If the ztC Edge system does not have Internet connectivity or if the **Upload** fails, you can manually upload the diagnostic file to the **Stratus Diagnostic Upload** web page. First, click **Download** on the ztC Console to download the diagnostic file as a .zip file to your local computer. Transfer the diagnostic zip file to a computer with Internet connectivity . Open a web browser, and in its address bar, enter http://diags.stratus.com/DiagUpload.html. On the **Stratus**

Diagnostic Upload page, click **Choose File**, select the zip file on the computer, and click **Submit**.

If you need help with this procedure, call ztC Edge Customer Support at the phone number listed on the **ztC Edge Support** page at https://www.stratus.com/services-support/customer-support/?tab=ztcedge.

After you are certain that you no longer need the file (for example, Customer Support confirms that the file uploaded correctly), you can optionally delete it from the ztC Edge system, as described in Deleting a Diagnostic File.

Related Topics

The ztC Console

The Preferences Page

Using the ztC Console

Deleting a Diagnostic File

Delete a diagnostic file from the ztC Edge system after you have uploaded it to your authorized Stratus service representative.

To delete a diagnostic file

1. Click **Support Logs** in the left-hand navigation panel, to open the **Support Logs** page.
2. Select the diagnostic file and click **Delete**.

Related Topics

The ztC Console

The Preferences Page

Using the ztC Console

The Physical Machines Page

The **Physical Machines** page enables you to manage the physical machines (PMs) in the ztC Edge system. (PMs are also referred to as nodes.) To open this page, click **Physical Machines** in the left-hand navigation panel.

State, **Activity**, **Name**, **Model**, and **# of VMs** columns appear immediately under the **PHYSICAL MACHINES** heading and masthead. To manage a specific PM, click **node0 (primary)** or **node1** (if it exists) under **Name**. To interpret PM states and activities, see Physical Machine States and Activities. To display information about a node, you can use the `snmptable` command; see Obtaining System Information with snmptable.

The bottom pane displays action buttons for and details about the selected node:

- Action buttons: Various action buttons appear, with inactive buttons grayed out, depending upon the state of the selected node. To perform most maintenance tasks, click **Work On**, which places a node into maintenance mode (for information, see Maintenance Mode). To learn about additional PM actions available in maintenance mode, see Physical Machine Actions or the help topic for the task you want to complete.
- Detailed information: To view detailed information or statistics about the selected node, click one of the following tabs:
  - **Summary** (in the initial display), which displays information about the node, such as (if applicable) the manufacturer, the model, serial number, overall state, activity, and configuration (memory and logical disks) for the selected node.
  - **Description**, which displays a window where you can enter information about the node.
  - **Storage**, which displays the state, logical ID, size, and size used of storage.
  - **Network**, which displays the state, name, speed, and MAC address of networks.
  - **Virtual Machines**, which displays the state, activity, and name of virtual machines.
  - **USB Devices**, which lists any USB devices inserted in the node.

Related Topics

The ztC Console

Using the ztC Console

Physical Machine Actions

When you select a physical machine (PM), some or all of the following action buttons appear, with inactive buttons grayed out, depending on the PM's state and activity.

> **Caution**: Use the **Physical Machines** page of the ztC Console when you perform maintenance on a PM. Avoid using controls on the computer, because the ztC Console protects the ztC Edge system from most actions that are potentially disruptive.

Commands	Description
 Work On	Enters a PM into maintenance mode. VMs running on this PM migrate to the other PM, if it exists and is in service. (Otherwise, you are asked to re-confirm the request and then shut down VMs.) When VMs are migrated or shut down, the PM displays **running (in Maintenance)**. See Maintenance Mode.
The following actions are available on some systems after clicking the **Work On** button, when the PM has entered maintenance mode.	
 Finalize	Removes a PM from the state **running (in Maintenance)**. See Maintenance Mode.
 Shutdown	Shuts down a PM. The PM transitions to **off (in Maintenance)**. See Shutting Down a Physical Machine.
 Reboot	Reboots the PM. The PM transitions to **preparing for reboot (in Maintenance)**. See Rebooting a Physical Machine.
 Remove	Causes the Stratus Redundant Linux software to delete the PM from the ztC Edge system's database, so that you can replace the PM or one of its components. See Replacing Physical Machines (Manual).
The following action may be available when a PM has failed or when the Stratus Redundant Linux software	

Commands	Description
has removed a PM from service and powered it off, due to an excessive failure rate.	
 Recover	Recovers a failed PM. In some cases, the ztC Console displays the state of a failed PM as **Unreachable (Syncing/Evacuating…)**. See Recovering a Failed Physical Machine (Manual).

Related Topics

The ztC Console

Using the ztC Console

The Physical Machines Page

Physical Machine States and Activities

The following states and activities apply to physical machines (PMs). Only certain actions are enabled during each state and activity.

State	Activity	Available Commands	Description
	 Running	**Work On**	PM is running normally.
	 Evacuating	**Finalize**	Virtual machines are migrating from this PM to its partner.
	 Running	**Work On**	PM is predicted to fail.
	 Running	**Work On**	PM failed.
	 Powered Off	**Work On**	ztC Edge has powered off the PM because of an excessive failure rate.
	 Booting	**Finalize**	PM is booting.

State	Activity	Available Commands	Description
	 Rebooting	**Finalize**	PM is rebooting.
	 Running	**Finalize** **Shutdown** **Reboot** **Recover** **Replace**	PM is running in Maintenance Mode. See Maintenance Mode.

Related Topics

The ztC Console

Using the ztC Console

The Physical Machines Page

The Virtual Machines Page

Use the **Virtual Machines** page to manage the virtual machines (VMs) running on your ztC Edge system. To open this page, click **Virtual Machines** in the left-hand navigation panel of the ztC Console.

To manage a specific VM, click the name of a VM in the top pane of the **Virtual Machines** page. The bottom pane displays controls and information for managing the VM.

To interpret VM status as displayed on the **Virtual Machines** page, see Virtual Machine States and Activities. To learn more about the controls on this page, see Virtual Machine Actions or the help topic for a specific task.

You can use the **Virtual Machines** page for administrative tasks including:

- Viewing information about a VM, including its name, operating system, description, and resources in the tabs of the bottom pane
- Creating, copying, exporting, importing, or restoring VMs, as described in Creating and Migrating Virtual Machines
- Opening a Virtual Machine Console Session
- Reprovisioning Virtual Machine Resources

- Controlling the power state of a VM, as described in:
  - Starting a Virtual Machine
  - Shutting Down a Virtual Machine
  - Powering Off a Virtual Machine
- Removing a Virtual Machine or Renaming a Virtual Machine
- Performing advanced tasks or troubleshooting, as summarized in Advanced Topics (Virtual Machines)
- Mounting (and unmounting) a USB device or a network-mounted folder for use by the guest operating system, as described in Mounting a USB Device or Network-mounted Folder on the ztC Edge System
- Attaching (and detaching) as USB device to a VM, as described in Attaching a USB Device to a Virtual Machine

Users who are assigned the role **Administrator** or **Platform Manager** can perform all tasks on the **Virtual Machines** page. Users who are assigned the role **VM Manager** can perform all tasks, except the **VM Manager** cannot expand a volume. For details on the **VM Manager** privileges, see Managing Virtual Machines. For information on assigning these roles, see Managing Local User Accounts.

Related Topics

Managing Virtual Machines

Using the ztC Console

Virtual Machine Actions

When you select a virtual machine (VM), the following action buttons can appear, with inactive buttons grayed out, depending on the VM's state and activity.

Action	Description
 Create	Launches the Create VM Wizard. See Creating a New Virtual Machine.
 Copy	Copies an existing VM on your system to create a new VM or to create a duplicate VM for troubleshooting. See Copying a Virtual Machine.

Action	Description
 Import/Restore	Imports a VM from a set of OVF and VHD files. See Creating and Migrating Virtual Machines. The import wizard allows you to *import* a VM to create a new instance of the VM or *restore* a VM to create an identical VM with the same hardware IDs provided in the OVF and VHD files. Open Virtual Machine Format (OVF) is an open standard for packaging and distributing physical or virtual machine data. The OVF format contains meta-data information about the VM. A Virtual Hard Disk (VHD) is a file that contains the virtual disk information.
The following actions are available for use if the VM is running.	
 Mount	Mounts a USB device or a network-mounted folder (that is, a directory) to make it available to the guest operating system. You can then export a VM to the mounted location. See Mounting a USB Device or Network-mounted Folder on the ztC Edge System.
 Unmount	Unmounts a mounted USB device or a network-mounted folder. See Mounting a USB Device or Network-mounted Folder on the ztC Edge System.
 Console	Opens a console for the selected VM. See Opening a Virtual Machine Console Session.
 Shutdown	Shuts down the selected VM. See Shutting Down a Virtual Machine.
 Power Off	Immediately stops processing in the selected VM and destroys its memory state. Use this only as a last resort, when the VM cannot be successfully shutdown. See Powering Off a Virtual Machine.

Action	Description
The following actions are available if the VM is shut down or stopped.	
 Config	Launches the **Reprovision Virtual Machine** wizard. The VM must be shut down prior to launching this wizard. See Reprovisioning Virtual Machine Resources.
 Restore	Recovers an existing VM on your ztC Edge system by overwriting the VM from a previous backup copy of OVF and VHD files. See Replacing/Restoring a Virtual Machine from an OVF File.
 Export	Saves the image of a VM to a set of OVF and VHD files. You can import these files on another system or import them back to the same ztC Edge system to restore or duplicate the original VM. See Exporting a Virtual Machine .
 Start	Boots the selected VM. See Starting a Virtual Machine.
 Boot From CD	Boots a VM from the selected virtual CD. See Booting from a Virtual CD.
 Remove	Removes a VM. See Removing a Virtual Machine.
The following action is available if the Stratus Redundant Linux software has removed the VM from service and powered it off because an excessive failure rate.	
 Reset Device	Resets the mean time between failures (MTBF) counter for a VM so it can be brought back into service. See Resetting MTBF for a Failed Virtual Machine. When a VM crashes, the Stratus Redundant Linux software automatically restarts it, unless it has fallen below its MTBF threshold. If the VM is below the MTBF

Action	Description
	threshold, the Stratus Redundant Linux software leaves it in the crashed state. If necessary, you can click **Reset Device** to restart the VM and reset the MTBF counter.

Related Topics

Managing the Operation of a Virtual Machine

The Virtual Machines Page

Using the ztC Console

Virtual Machine States and Activities

A virtual machine (VM) can have the following states and activities, during which only certain actions are enabled.

State	Activity	Enabled Actions	Description
	 Installing		The Stratus Redundant Linux software is installing the boot volume for a new VM.
	 stopped	Start Copy Config Export Boot From CD Remove	VM has been shut down or powered off.
	 booting	Console Power Off	VM is starting.

State	Activity	Enabled Actions	Description
	 running	Console Shutdown Power Off	VM is operating normally on redundant physical machines
	 running	Console Shutdown Power Off	VM is operating normally, but is not running on fully redundant resources.
	 stopping	Power Off Remove	VM is being shut down in response to the **Shut-down** action, or shut down because the remaining physical machine is transitioning into maintenance mode.
	 crashed		VM crashed and is restarting. If enabled, e-Alerts and support notification messages are sent.
	 crashed		VM crashed too many times and exceeded its MTBF threshold. The VM is left in a crashed state until **Reset Device** is clicked. See Resetting MTBF for a Failed Virtual Machine.

Related Topics

Managing the Operation of a Virtual Machine

The Virtual Machines Page

Using the ztC Console

The Volumes Page

The **Volumes** page displays information about volumes that are attached to virtual machines (VMs) in the ztC Edge system. To open this page, click **Volumes** in the left-hand navigation panel of the ztC Console. The **Volumes** page displays the following columns with information about volumes in the top pane:

- **State**
- **Name**
- **Disk Synchronization** (if it exists)
- **Size**
- **Bootable**
- **Used By**, which displays one of the following:
  - A link to a VM when a VM is using the volume.
  - A link to the physical machine (PM) page (**node0** or **node1**, if it exists) when the volume is **root** or **swap**.
  - ***System*** for a shared volume (**shared.fs**) .
  - ***None*** when the volume is not a system volume and is not used by a VM.

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. (You can also display information about volumes using the `snmptable` command; see Obtaining System Information with snmptable.) You can perform some administrative tasks on volumes from the bottom pane, including:

- Add a description for each volume in the **Description** text box.
- Rename a volume (see Renaming a Volume on the ztC Edge System).
- Remove a volume by clicking **Remove**. Note, though, that the **Remove** button is grayed out when a VM is using a volume.

You perform other volume management tasks from the virtual machines page. These tasks include:

- Attaching a Volume to a Virtual Machine
- Creating a Volume in a Virtual Machine
- Detaching a Volume from a Virtual Machine
- Removing a Volume from a Virtual Machine

Related Topics

The ztC Console

Using the ztC Console

The Networks Page

The **Networks** page displays information about the shared networks attached to the ztC Edge system. To open this page, click **Networks** in the left-hand navigation panel of the ztC Console.

You can use the **Networks** page to view information about a specific network, including its state, link condition, name, internal name, type, number of connected Virtual Machines (VMs), speed, and MTU. You can also add a description for the network using the **Description** tab in the bottom pane.

To manage a specific network or simply view information about it, click the network name under **Name** or **Internal Name** in the top pane of the **Networks** page, or click a port in the network connectivity diagram on the **Summary** tab. The bottom pane displays additional information about nodes on the network. Columns in the **Summary** tab display information about the node's state, physical interface, speed, MAC address, slot, and port. To display or hide columns, move the cursor to the right of a column heading, click the down-arrow that appears, and then click **Columns**, selecting or de-selecting the columns that you want to show or hide.

You can use the **Networks** page for administrative tasks, including:

- Viewing a list of the physical adapters that compose the network, on the **Summary** tab.
- Adding a description for a network, on the **Description** tab.
- Viewing a list of virtual machines that use the network, on the **Virtual Machines** tab.
- Changing the name by double-clicking the name in the **Name** column.
- Setting the MTU for A-Link and business networks.

For additional information on networks, see the following:

- Network Architecture
- Connecting Ethernet Cables
- General Network Requirements and Configurations
- Meeting Network Requirements for ALSR configurations

Note: The **Networks** page displays only networks that have physical connectivity on both physical machines. If a network that you expect to see does not appear, check that both network connections are cabled correctly and that their LINK is active.

Related Topics

The ztC Console

Using the ztC Console

Setting the MTU

Network performance improves with the highest maximum transmission unit (MTU) that the network can support. You can specify the MTU value for A-link and business (biz) networks using the **Networks** page of the ztC Console.

> **Note**: If you change the MTU of either business network `ibiz0` (also referred to as `network0`) or `ibiz1` (also referred to as `network1`), the system automatically migrates the VMs from one node to the other. If you change the MTU for `ibiz0` specifically, the system also automatically fails over from the primary node to the secondary node. To avoid this issue, avoid changing the MTU of the business networks, or change the MTU only during a planned maintenance period.

To set the MTU of an A-Link or business network

1. Click **Networks** in the left-hand navigation panel, to open the **Networks** page.
2. In the top pane, select the A-link or business network whose MTU you want to set.
3. Click **Config**.
4. In the **Configure Shared Network** window, select the **Network Role** (**Business** or **A-Link**).
5. Under **MTU**, type a bytes value from 1280 to 65535 (the default is 1500).
6. Click **Save**.

Related Topics

The Networks Page

The ztC Console

Using the ztC Console

The Virtual CDs Page

Use the **Virtual CDs** page to create virtual CDs (VCDs). Use VCDs to make software installation or recovery media available to the virtual machines on the system. To open this page, click **Virtual CDs** in the left-hand

navigation panel of the ztC Console.

To manage a specific VCD, click the name of a VCD in the top pane of the **Virtual CDs** page. The bottom pane displays a description of the VCD.

You can use the **Virtual CDs** page for administrative tasks including:

- Creating a Virtual CD
- Removing a Virtual CD
- Renaming a Virtual CD
- Adding a description for each volume, in the **Description** text box

To complete other VCD management tasks, see Managing Virtual CDs.

Related Topics

Inserting a Virtual CD

Ejecting a Virtual CD

Using the ztC Console

The Upgrade Kits Page

The **ztC Edge Upgrade Kits** page enables you to upload and manage upgrade kits that you use to upgrade the system to newer versions of the Stratus Redundant Linux software. You can also copy an upgrade kit to a USB medium in order to use the medium when reinstalling the system software.

To open the **Upgrade Kits** page, click **Upgrade Kits** in the left-hand navigation panel in the ztC Console.

For information about upgrading the Stratus Redundant Linux software, see Upgrading Stratus Redundant Linux Software.

For information about creating a USB medium, see Creating a USB Medium with System Software.

Related Topics

The ztC Console

Using the ztC Console

Creating a USB Medium with System Software

You can use the **Upgrade Kits** page to create a USB medium with a copy of the deployment ISO file of the system software, Stratus Redundant Linux. You then use the USB medium to reinstall the software if you

need to manually recover or replace a failed node.

Note: Copying an upgrade kit to a USB medium dismounts file systems, if any, from the medium.

To create a USB medium with system software

1. Download an upgrade kit, if you have not already done so. See Upgrading Stratus Redundant Linux Software Using an Upgrade Kit.
2. Insert a USB medium into the primary node. On the **Physical Machines** page, check that the **USB Devices** tab lists the device.
3. In the ztC Console, click **Upgrade Kits** in the left-hand navigation panel.
4. If the **Upgrade Kits** page lists more than one kit, select the version with the ISO that you want to copy.
5. Click the **Create USB Medium** button (beneath the masthead).

   The **Create USB Medium** dialog box opens.
6. If the node has more than one USB medium, you need to select a medium from the drop-down list. Then, click **Continue** (or click **Cancel** to cancel the procedure).

   The **Create USB Medium** dialog box displays the percentage of progress. The window closes when copying has finished.

Use the USB medium to reinstall the software if you need to manually recover or replace a failed node. See Recovering a Failed Physical Machine (Manual) or Replacing Physical Machines (Manual).

Related Topics

The Upgrade Kits Page

The Users & Groups Page

Use the **Users & Groups** page to add, modify, or remove user accounts on the ztC Edge system.

To open this page, click **Users & Groups** in the left-hand navigation panel of the ztC Console.

To manage local user accounts

To add a new user, click **Add** in the right side of the top pane. To modify an existing user, click the name of a user account and click **Edit** or **Remove**.

For more information, see Managing Local User Accounts.

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**, **Real Name**, **Email** address, or **Role**.

Related Topics

Managing Local User Accounts

Managing Local User Accounts

You add, edit, or remove users, specify passwords, and assign user roles to local-user accounts on the **User & Groups** page in the ztC Console.

Local user accounts reside on the ztC Edge system itself, as opposed to a central domain server. You can find local accounts on the **Users & Groups** page by looking for entries labeled **Local User** in the **Type** column.

User roles are:

- **Administrator**: full system administrator privileges
- **Platform Manager**: system administrator privileges except for adding, removing, and modifying users
- **VM Manager**: ability to manage VMs (see Managing Virtual Machines for detailed information)
- **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 & Groups** page: in the left-hand navigation panel, select **Users & Groups**.

To add a user account

1. In the top pane, click **Add**.
2. In the **Role** drop-down window, select **Administrator**, **Platform Manager**, **VM Manager**, or **Read-only**.
3. Provide values for the **User Name**, **Password** (and **Confirm Password**), **Email Address**, and **Real Name** fields. User names and passwords may be from 1 to 64 characters long, and must include no white space.
4. Click **Save**.

To edit a user account

1. Select the account you want to edit.

2. In the top 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**, **VM Manager**, or **Read-only**.
4. Click **Save**.

To remove a user account

1. Select the account to remove.
2. Click **Remove** in the top pane.
3. Click **Yes** in the Confirm dialog box.

Note: You cannot delete the default **admin** account, although you should change its name and password by editing the account.

Related Topics

The Users & Groups Page

4

Chapter 4: Upgrading Stratus Redundant Linux Software

To upgrade Stratus Redundant Linux software, use an upgrade kit. See Upgrading Stratus Redundant Linux Software Using an Upgrade Kit.

Related Topics

The Upgrade Kits Page

The ztC Console

Using the ztC Console

Upgrading Stratus Redundant Linux Software Using an Upgrade Kit

This topic describes how to use an upgrade kit of Stratus Redundant Linux software to upgrade the system software. The topic also explains how to download the kit and then upload it to the system, if you need to do so before upgrading the system. For systems licensed for two nodes, you can optionally control the upgrade by enabling pauses (you cannot enable pauses on a system licensed for one node). Inspecting a system during a pause is useful for verifying or reconfiguring third-party tools or other services that are not managed by the system.

Caution: Do not update the CentOS host operating system of the ztC Edge system from any source other than Stratus. Use only the CentOS release that is installed with the Stratus Redundant Linux software.

Prerequisites:

- All PMs and VMs must be in good health before upgrading the system software. Before starting an upgrade, examine the ztC Console to verify that there are no alerts indicating PM or VM problems.
- Eject any VCDs or USB media from the VMs before upgrading the system software. If VCD or USB media is still connected to the VMs, it prevents the system from migrating the VMs and putting the PMs into maintenance mode for the upgrade process.
- To verify that the system meets the requirements of the upgrade kit, use the **Qualify** button, as described in this topic.
- Before you upgrade a system licensed for one node, you should back up the VMs. Then, upgrade and qualify the software following the instructions below. Finally, upgrade the one PM of the system using the procedure in To upgrade a system licensed for one node. The upgrade includes at least a 15-minute loss of access to the ztC Console as the system reboots during the upgrade procedure.

The steps are:

I. To download the upgrade kit

II. To upload the upgrade kit to the system

III. To qualify the software (optional)

IV. To upgrade the system software

I. To download the upgrade kit

Note: The **Upgrade Kits** page of the ztC Console allows only two saved kits. If the pages lists two kits and you want to download another kit, you first need to delete a kit.

1. Open the **Downloads** page at https://www.stratus.com/services-support/downloads/?tab=ztcedge.
2. Scroll down to the upgrade section and then click the upgrade link to download the kit.
3. Navigate to a location on a local computer to save the file. If necessary, transfer the file to the remote management computer running the ztC Console.

II. To upload the upgrade kit to the system

Upload the upgrade kit, if necessary, to the ztC Edge system from the remote management computer that is running the ztC Console.

1. In the ztC Console, click **Upgrade Kits** in the left-hand navigation panel.
2. On the **Upgrade Kits** page, click the **Add a Kit** button beneath the masthead, which opens the **ztC Edge - Kit Upload Wizard**.
3. In the **ztC Edge - Kit Upload Wizard** dialog box, click **Choose File** (in Google Chrome) or **Browse** (in Firefox or Internet Explorer), and then browse to select a .kit file.
4. After you have selected a .kit file, click **Upload**, **Import**, or **Finish** (they perform the same function). A message such as **Uploading file (DO NOT CLOSE WIZARD)** appears while the file is uploading. The upload may require up to two minutes for a file stored locally, to ten or more minutes for a file stored over a network. If the upload fails, the wizard displays the message **Failed to upload file**.
5. After the upload is complete, the wizard closes and the **Upgrade Kits** page lists the state and version number of the upgrade kit. The **Qualify**, **Upgrade**, and **Delete** buttons also appear with the **Add a Kit** button.
6. If more than one upgrade kit is loaded, select the one to use.

III. To qualify the software

Qualify the software to verify that your system meets the requirements of the upgrade kit. (Qualifying the software is recommended, but not required.)

To do so, select the upgrade kit you want to qualify on the **Upgrade Kits** page, and then click **Qualify**.

The qualification may require up to six minutes. If the qualification succeeds, continue with the next step.

If the qualification fails, a pop-up window appears with messages indicating the cause of the failure. These messages may indicate unsupported releases, insufficient storage, partition problems, VMs that need to be shutdown, or other information associated with upgrading the system. For example, if the system has insufficient disk space to complete the upgrade, the message `Insufficient free space` appears reporting the amount of space needed. If you need help resolving a qualification issue, search for the qualification error message in the **Knowledge Base** in the **Stratus Customer Service Portal** at https://support.stratus.com.

IV. To upgrade the system software

1. Begin the upgrade by clicking **Upgrade** on the **Upgrade Kits** page.

   A **Confirm** window appears, stating that you have chosen to upgrade the system and displaying a message asking you to confirm the upgrade to the selected upgrade kit. The window also includes a check box for you to enable pauses, allowing you to control the upgrade. Enable pauses by clicking the box **Pause after individual node upgrades**.

2. Click **Yes** to continue the upgrade.

   The upgrade begins. If you enabled pauses, the diagram illustrating the upgrade steps displays the current state of the upgrade. When the upgrade pauses, you must click **Finalize** to continue.

After one node has been upgraded, but the other node (if it exists) has not yet been upgraded, the nodes are running different versions of the software. During this time, the masthead displays the message **System is running with mismatched versions**.

After the upgrade is complete, check for updated virtIO drivers on all Windows-based VMs, as described in Updating the VirtIO Drivers (Windows-based VMs).

To upgrade a system licensed for one node

1. Shut down all VMs that are running on the ztC Edge system.
2. Upgrade the system with an upgrade kit, using the instructions in the steps above.

   Note: The upgrade includes at least a 15-minute loss of access to the ztC Console as the system reboots during the upgrade procedure.

3. Ensure that the system is running correctly.
4. Start all of the VMs.

Related Topics

The Upgrade Kits Page

The ztC Console

Using the ztC Console

ztC Edge System Description

5

Chapter 5: Managing Physical Machines

Manage a physical machine (PM), or node, to control its operation and perform maintenance.

You view and manage PMs using the **Physical Machines** page of ztC Console; for information, see The Physical Machines Page.

Many of the tasks you perform from the **Physical Machines** page require maintenance mode; for information, see Maintenance Mode.

To manage the operational state of a PM (in maintenance mode), see:

- Rebooting a Physical Machine
- Shutting Down a Physical Machine
- Load Balancing

To power on a PM (at the physical console of the PM), see Powering On a Physical Machine.

To troubleshoot a PM by recovering a failed PM or resetting MTBF for a failed machine, see Troubleshooting Physical Machines.

To perform maintenance tasks on the PM hardware, such as replacing a PM, see Maintaining Physical Machines.

Maintenance Mode

When a physical machine (PM) enters maintenance mode, it goes offline for service. When you finalize service, the PM exits maintenance mode and goes back online, becoming available for running virtual machines (VMs).

For a system licensed for two nodes (that is, two PMs), note the following:

- When one PM enters maintenance mode, the PM migrates the VMs running on it to the other PM, which protects the VMs from any potential disruption caused by the service. When both PMs enter maintenance mode, the PMs perform an orderly shutdown of all VMs, which protects their memory state before the PMs shut down or reboot.
- When the primary PM (**node*x* (primary)**) enters maintenance mode, the other PM becomes primary.
- If you want both PMs in maintenance mode, first enter the secondary PM into maintenance mode, and then enter the primary PM into maintenance mode. This sequence avoids the unnecessary migration of VMs.

For a system licensed for one node (that is, one PM), the PM shuts down VMs when it enters maintenance mode. So, place the PM in maintenance mode only during a planned maintenance period.

Shut down the PMs only from the **Physical Machines** page with the PM in maintenance mode because the ztC Console protects the system from disruptive action that results from manually powering down a PM.

Cautions:

1. The system is not redundant when a PM is in maintenance mode. For continuous uptime, finalize service as soon as possible so that the PM can exit maintenance mode and go back online.
2. Place all PMs in maintenance mode only if you are able to shut down all business processing. If you need to keep VMs running on a system licensed for two PMs, avoid entering both PMs into maintenance mode at the same time. To keep VMs running, at least one PM must be up and running normally. (If you need to shut down the entire ztC Edge system, see Shutting Down a Physical Machine.)

To enter a PM into maintenance mode

1. Select a PM from the **Physical Machines** page.
2. Click **Work On**.

   When the PM is in maintenance mode, its state displays  .

To finalize and exit a PM from maintenance mode

1. Select a PM from the **Physical Machines** page.
2. Click **Finalize**, which exits the PM from maintenance mode.

Related Topics

The ztC Console

Managing Physical Machines

Physical Machines and Virtual Machines

The Physical Machines Page

The Virtual Machines Page

Powering On a Physical Machine

Power on a physical machine (PM) at the physical console of the PM.

> **Note**: If a PM loses power because you disconnect the power cord or the AC mains power is lost, each PM in a ztC Edge system is set to power on automatically as soon as power is restored.

To power on a PM

1. Press the power button on the front panel of the PM.
2. Ensure that the **PWR** LED on the front panel is lit.

If you want to power on the system, press the power button on the front panel of each PM in the system, as described in Powering On the System.

Related Topics

Maintenance Mode

The ztC Console

Managing Physical Machines

The Physical Machines Page

Rebooting a Physical Machine

Reboot a physical machine (PM) to restart its Stratus Redundant Linux software, and optionally exit the PM from maintenance mode. (If you need to reboot both PMs in a system licensed for two nodes, see Rebooting the System.)

If you are rebooting the one PM of a single-node system, do so only during a planned maintenance period, since rebooting the PM will shutdown VMs and you need to manually restart them.

To reboot a PM

1. Determine which PM (node0 or node1, if it exists) you want to reboot.
2. In the ztC Console, click **Physical Machines** in the left-hand navigation panel.
3. Select the appropriate PM (node0 or node1, if it exists) and then click **Work On**, which changes the PM's **Overall State** to **Maintenance Mode** and the **Activity** state to **running (in Maintenance)**.
4. Click **Reboot**. As the PM reboots, the **Activity** state displays:
   - **preparing for reboot (in Maintenance)**
   - **rebooting (in Maintenance)**
   - **booting (in Maintenance)**
   - **running (in Maintenance)**.
5. To exit the PM from maintenance mode and make it available for running virtual machines, click **Finalize**.

   On a single-node system, you need to manually restart VMs.

Related Topics

Maintenance Mode

The ztC Console

Managing Physical Machines

The Physical Machines Page

Shutting Down a Physical Machine

Shut down a physical machine (PM), or node, to stop it from running when you need to service or replace the PM. Use the following procedures to shut down one and only one PM by using the ztC Console or the power button on the PM.

Cautions:

1. Data loss will occur if you use the following procedures to shut down both PMs of a ztC Edge system licensed for two nodes or the one PM of a system licensed for one node. If you need to stop both PMs of a system licensed for two nodes or the one PM of a system licensed for one node, shut down the system (which also shuts down the virtual machines (VMs)), as described in Shutting Down the System.
2. Do not use the `-f` (force) option with the `halt`, `poweroff`, or `reboot` command of the host operating system of a PM. Doing so causes guests that are active on the same PM to hang.
3. The ztC Edge system is not redundant when you shut down a PM. Bring an offline PM back into service as soon as possible.

Note: When you shut down a PM, standby power remains on for lights-out management unless you disconnect the power cord or AC mains power is lost.

To shut down a PM in the ztC Console

To shut down a PM, you must place the PM into maintenance mode, which migrates any VMs running on that PM to the remaining PM (if it exists). On a system with two nodes, the VMs continue to run during this process, which takes a minute or two.

1. Determine which PM you want to shut down.
2. In the ztC Console, click **Physical Machines** in the left-hand navigation panel.
3. Select the appropriate PM (node0 or node1, if it exists) and then click **Work On**, which changes the PM's **Overall State** to **Maintenance Mode** and the **Activity** state to **running (in Maintenance)**.
4. After the PM displays **running (in Maintenance)**, click **Shutdown**.

After the PM has shut down, its activity is **off (in Maintenance)**. You must manually restart the PM.

To shut down a PM with the power button

To shut down a PM by using the power button

You can also shut down a PM by using the power button on the PM. On a system licensed for two nodes, the VMs continue to run during this process, which takes a minute or two.

1. While both PMs are running, press and release the power button on the front panel of the PM that you want to shut down.
2. The system automatically places the PM into maintenance mode, which migrates the VMs running on the PM to the remaining PM.
3. The PM automatically shuts down.

When the PM shuts down, the **PWR** LED on the front panel is off, though standby power remains. You must manually restart the PM.

To forcibly power off a PM by using the power button

> **Caution**: If a PM does not turn off after you click **Shutdown** or press the power button, you may need to forcibly power off the PM. Forcibly powering off a PM destroys its memory state; therefore, do this only as a last resort.

Press the power button on the PM for several seconds to forcibly remove power.

When the PM powers off, the **PWR** LED on the front panel is off, though standby power remains. You must manually restart the PM.

Related Topics

Maintenance Mode

The ztC Console

Managing Physical Machines

The Physical Machines Page

Load Balancing

On systems licensed for two nodes, HA Load Balancing distributes VMs across both PMs to improve performance and availability. Load balancing is configured per VM and is enabled automatically on ztC Edge systems. (Systems licensed for one node do not provide load balancing.)

If a PM is out of service, all VMs will run on the surviving PM. VMs automatically migrate back as soon as the PM they are targeted to run on returns to service and is fully synchronized.

Modes of Operation

Load balancing is set for a VM on its **Load Balance** tab on the **Virtual Machines** page. The following modes are supported:

- **automatically balance**. This provides automatic load balancing of a VM. When a VM is set to balance automatically, it will run on the available PM with the most resources. When the system determines that better load balancing can be achieved by moving one or more VMs with the automatic setting, an alert is generated. The alert appears on the Dashboard, and a Load Balancing notification appears on the masthead. In response to the alert, click **Load Balance** in the masthead to initiate automatic load balancing of a VM.
  The icon on the **Virtual Machines** page under **Current PM** column indicates VMs that will migrate imminently.
- **manually place on node*N***. Advanced users can manually assign a preferred PM (node) for each individual VM, rather than relying on the automatic policy, if preferred.

A graphic appears on the **Virtual Machine** page in the **Current PM** column for each VM. It indicates the current status of the VM's load-balancing state, the PM the VM is running on, and its preference.

The following sample graphic indicates that the VM is currently on PM 0 and that PM 1 is the preference.

ztC Edge policy ensures that a VM is always running. In the event that one PM is predicted to fail, is under maintenance, or is taken out of service, the VM will run on the healthy PM. When both PMs are healthy, a VM migrates to its preferred PM.

Related Topic

Selecting a Preferred PM for a Virtual Machine

Troubleshooting Physical Machines

The following topic describes troubleshooting procedures for PMs:

- Recovering a Failed Physical Machine (Manual)

If you cannot recover a PM using the software-based troubleshooting procedure above, see Maintaining Physical Machines for information about replacing the physical PM hardware.

Recovering a Failed Physical Machine (Manual)

Caution: If you need to recover or replace a PM in a ztC Edge system, use the instructions in *ztC Edge 100i Systems: Node Replacement* (R002Z). (If needed, see Replacing Physical Machines (Automated) for additional details.) Avoid using the manual procedure described in this topic unless specifically instructed by your authorized Stratus service representative.

Recover a physical machine (PM), or node, when it cannot boot or if it fails to become a PM in the ztC Edge system. In some cases, the ztC Console displays the state of a failed PM as **Unreachable (Syncing/Evacuating)**.

To recover a PM, you must reinstall the Stratus Redundant Linux release that the PM has been running. Recovering a failed PM, though, is different from installing the software for the first time. The recovery preserves all data, but it re-creates the /boot and root file systems, re-installs the Stratus Redundant Linux system software, and attempts to connect to the existing system. (If you need to replace the physical PM hardware instead of recovering the system software, see Replacing Physical Machines (Manual).)

To reinstall the system software, you can allow the system to automatically boot the replacement node from a temporary Preboot Execution Environment (PXE) server on the primary PM. As long as each PM contains a full copy of the most recently installed software kit (as displayed on the **Upgrade Kits** page of the ztC Console), either PM can initiate the recovery of its partner PM with PXE boot installation. If needed, you can also manually boot the replacement node from USB installation media.

Use one of the following procedures based on the media you want to use for the installation, either **PXE** or **USB** installation.

Caution: The recovery procedure deletes any software installed in the host operating system of the PM and all PM configuration information entered before the recovery. After you complete this procedure, you must manually re-install all of your host-level software and reconfigure the PM to match your original settings.

Prerequisites:

1. Determine which PM you need to recover.
2. If you want to use a USB medium to install the system software on the replacement PM, create a bootable USB medium as described in Creating a USB Medium with System Software.

   When creating the USB medium, ensure that it contains the most recently installed upgrade kit. For example, if the release shown in the masthead of the ztC Console window is version `1.2.0-550`, where 550 is the build number, the kit you select to create the USB medium on the **Upgrade Kits** page must also be version `1.2.0-550`. If the system detects a different build on the target PM, it automatically overrides the recovery process, **initializes all data** on the target PM, and uses PXE boot installation to reinstall the most recently installed software kit on the PM with no user interaction.
3. If using a USB medium, connect a keyboard and monitor to the replacement PM to monitor the installation process and specify settings.

To recover a PM (with PXE boot installation)

Use the following procedure to recover a PM by using PXE boot installation to reinstall the system software from the software kit on the primary PM.

1. In the ztC Console, click **Physical Machines** in the left-hand navigation panel.
2. Select the appropriate PM (node0 or node1) and then click **Work On**, which changes the PM's **Overall State** to **Maintenance Mode** and the **Activity** state to **running (in Maintenance)**.
3. After the PM displays **running (in Maintenance)**, click **Recover**.
4. When prompted to select the type of repair, click **PXE PM Recover - Preserve Data**.

   Caution: It is important to select **PXE PM Recover: Preserve data**; otherwise, the installation process may delete data on the target PM.
5. Click **Continue** to begin the recovery process. The system reboots the target PM in preparation for the system software reinstallation.
6. The recovery process continues with no user interaction, as follows:

- The target PM begins to boot from a PXE server that temporarily runs on the primary node.
- The target PM automatically starts the system software installation, which runs from a copy of the installation kit on the primary node.
- The installation process reinstalls the system software, while preserving all data.

You do not need to monitor the progress of the software installation or respond to prompts at the physical console of the target PM. The recovery process is automated, and it is normal for the PM to display a blank screen for a long period of time during the software installation.

7. When the software installation is complete, the target PM reboots from the newly installed system software.
8. As the target PM boots, you can view its activity on the **Physical Machines** page of the ztC Console. The **Activity** column displays the PM as **(in Maintenance)** after the recovery is complete.
9. If applicable, manually reinstall applications and any other host-level software, and reconfigure the PM to match your original settings.
10. When you are ready to bring the target PM online, click **Finalize** to exit maintenance mode. Verify that both PMs return to the **running** state and that the PMs finish synchronizing.

> **Note**: When the target PM exits maintenance mode, the system automatically disables the PXE server on the primary node that was used for the recovery process.

To recover a PM (with USB installation)

Use the following procedure to recover a PM by reinstalling the system software from a USB medium.

1. In the ztC Console, click **Physical Machines** in the left-hand navigation panel.
2. Select the appropriate PM (node0 or node1) and then click **Work On**, which changes the PM's **Overall State** to **Maintenance Mode** and the **Activity** state to **running (in Maintenance)**.
3. After the PM displays **running (in Maintenance)**, click **Recover**.
4. When prompted to select the type of repair, click **USB PM Recover - Preserve Data**.

> **Caution**: It is important to select **USB PM Recover: Preserve data**; otherwise, the installation process may delete data on the target PM.

5. Click **Continue** to begin the recovery process. The system shuts down the target PM in preparation for the system software reinstallation.
6. Connect the bootable USB medium to the target PM, and then manually power on the PM.
7. As the target PM powers on, enter the firmware (UEFI) setup utility. In the **Save & Exit** menu, under **Boot Override**, select the **UEFI** entry for the USB medium to boot from the device one time during the next boot sequence. The PM restarts.

   Note: Use the **Boot Override** property to temporarily change the boot device instead of modifying the persistent **BOOT ORDER Priorities** in the **Boot** menu. The top boot priority must remain **UEFI Network** (default) to support the automated node replacement that is typically performed on ztC Edge systems.

8. Monitor the installation process at the physical console of the target PM.
9. At the **Welcome** screen, use the arrow keys to select the country keyboard map for the installation.
10. At the **Install or Recovery** screen, select **Recover PM, Join system: Preserving data** and press **Enter**. The recovery process continues with no user interaction.

    Caution: It is important to select **Recover PM, Join system: Preserving data**; otherwise, the installation process may delete data on the target PM.

11. When the software installation is complete, the target PM reboots from the newly installed system software.
12. As the target PM boots, you can view its activity on the **Physical Machines** page of the ztC Console. The **Activity** column displays the PM as **(in Maintenance)** after the recovery is complete.
13. If applicable, manually reinstall applications and any other host-level software, and reconfigure the PM to match your original settings.
14. When you are ready to bring the target PM online, click **Finalize** to exit maintenance mode. Verify that both PMs return to the **running** state and that the PMs finish synchronizing.

Related Topics

Maintenance Mode

Managing Physical Machines

The ztC Console

The Physical Machines Page

6

Chapter 6: Managing Virtual Machines

Manage a virtual machine (VM) to control its operation, provision its resources, or configure its guest operating systems and applications.

You can view and manage VMs on the **Virtual Machines** page of the ztC Console, which you access as described in The Virtual Machines Page. To perform specific management tasks, see the following topics.

To manage the operational state of a VM, see:

- Starting a Virtual Machine
- Shutting Down a Virtual Machine
- Powering Off a Virtual Machine
- Opening a Virtual Machine Console Session
- Renaming a Virtual Machine
- Removing a Virtual Machine

To display information about a VM, use the `snmptable` command (see Obtaining System Information with snmptable).

To create or configure a VM, see:

- Planning Virtual Machine Resources (virtual CPUs, memory, storage, and networks)
- Creating and Migrating Virtual Machines
- Managing Virtual CDs
- Configuring Windows-based Virtual Machines

- Configuring Linux-based Virtual Machines
- Managing Virtual Machine Resources

To attach a USB device to a VM, see Attaching a USB Device to a Virtual Machine.

To complete advanced tasks, see:

- Assigning a Specific MAC Address to a Virtual Machine
- Selecting a Preferred PM for a Virtual Machine
- Configuring the Boot Sequence for Virtual Machines
- Resetting MTBF for a Failed Virtual Machine

The local-user role of **VM Manager** can perform many of these tasks. Specifically, the **VM Manager** can:

- Perform tasks of the available function buttons and actions on The Virtual Machines Page.
- View all available tabs on The Virtual Machines Page.
- Create and delete VCDs from the The Virtual CDs Page.

For information on assigning the role of **VM Manager**, see Managing Local User Accounts.

Planning Virtual Machine Resources

When creating virtual machines, plan to allocate system resources to maximize system performance and continuous uptime.

To plan for allocating resources to your virtual machines, see:

- Planning Virtual Machine vCPUs
- Planning Virtual Machine Memory
- Planning Virtual Machine Storage
- Planning Virtual Machine Networks

Planning Virtual Machine vCPUs

Allocate virtual CPUs (vCPUs) to assign computing resources to a virtual machine (VM) on the ztC Edge system.

When allocating vCPUs to a VM, consider the following information and restrictions:

- Each vCPU represents a virtual unit of processing power. The total number of vCPUs available on a system is equal to the minimum of the number of hardware threads presented by either physical

machine (PM) in the system. For example, in a system where one PM that has 4 cores and 2 threads per core (8 vCPUs) and a second PM (in that system) that has 8 cores and 2 threads per core (16 vCPUs), the total number of vCPUs available is 8 vCPUs (fewest threads of either PM).

- The number of vCPUs available for the VMs is equal to the total number of vCPUs on the system.
- The maximum number of vCPUs you can allocate to any one VM is the total number of vCPUs in the system.
- Windows-based VMs: If you change the number of assigned vCPUs from 1 to *n* or *n* to 1, after restarting the VM at the end of the reprovisioning process (see Reprovisioning Virtual Machine Resources), you must shut down and restart the VM a second time. This allows the VM to correctly reconfigure itself for Symmetric Multiprocessing (SMP). The VM displays odd behavior and is not usable until it is restarted.
- The **System** page of the ztC Console (see The System Page) indicates the total number of vCPUs, the number of vCPUs allocated to the ztC Edge system software, the number of vCPUs consumed by running VMs, and the number of free vCPUs.
- The Stratus Redundant Linux software allows the over-provisioning of vCPUs. If the number of free vCPUs on the **System** page is less than zero, you have over-provisioned the vCPUs; the console indicates this and displays an estimate of the degree to which vCPUs have been over-provisioned.
- The over-provisioning of vCPUs does not prevent you from creating or starting VMs; however, it is best to avoid running the system in an over-provisioned state.

Considerations When Over-Provisioning Virtual CPUs

Note: In general, avoid over-provisioning VM resources. It is best to isolate each VM's resources to protect the VM against other VMs that might experience resource leaks or unexpected performance peaks. When you create and configure VMs, assign dedicated resources that cannot be used by other VMs.

You should over-provision physical CPUs only under the following conditions:

- The peak vCPU resources consumed by the combined VMs does not exceed the physical resources of the ztC Edge system.
- One or more VMs are used at different times (such as off-peak backups).

- One or more of the VMs will be stopped while the other is running, for example, during VM upgrades or VM point-in-time backup or recovery.
- Peak total CPU use by VMs will not affect service level agreements or required response times.
- Each VM's CPU use is well understood, and its application(s) are not prone to resource leaks. When CPUs are over-provisioned, a leak in one VM can affect the performance of other VMs.

Related Topics

System Requirements Overview

Creating and Migrating Virtual Machines

Managing Virtual Machine Resources

Planning Virtual Machine Memory

Allocate memory to assign physical memory to a virtual machine (VM) on your ztC Edge system.

When allocating memory to a VM, consider the following information and restrictions:

- The total memory you can allocate to the VMs is equal to the total amount of memory available on the ztC Edge system (see System Requirements Overview) minus the memory allocated to the ztC Edge system software. For example, if the total amount of memory is 32 GB, and 2 GB is allocated to the system software, there are 30 GB of memory available to the VMs.
- For systems licensed for two nodes, you can provision a single VM with memory up to the total amount of memory available to the VMs. Each VM consumes its requested amount of memory plus an additional 20% memory for overhead.
- The minimum memory allocation is 256 MB, but 64-bit operating systems require 600 MB or more. Be sure to verify the memory requirements for your guest operating systems.
- The **System** page of the ztC Console (see The System Page) indicates the total amount of memory, the memory allocated to the ztC Edge system software, the memory consumed by running VMs, and the amount of free memory. Use this page to verify your memory allocations.
- The Stratus Redundant Linux software does not allow over-provisioning of memory for **running** VMs; it prevents you from starting VMs that would exceed the total physical memory of the physical machines. You may safely allow over-provisioning of memory to occur only if one or more of the VMs is **stopped** while the other is running, for example, during VM upgrades or VM point-in-time backup or recovery.

- If necessary, you can manually redistribute memory by shutting down or reconfiguring one or more under-utilized VMs and then reassigning the available resources to a more heavily-utilized VM.

Related Topics

System Requirements Overview

Creating and Migrating Virtual Machines

Managing Virtual Machine Resources

Planning Virtual Machine Storage

Plan the allocation of storage on your ztC Edge system to ensure that you have space for your virtual machines (VMs) and system management needs.

When allocating storage to your virtual machines (VMs), consider the following actions:

- Observe storage maximums

  The Stratus Redundant Linux software does not allow over-provisioning of storage. The aggregate required storage for all VMs and VCDs must be no more than the total available storage in the ztC Edge system.

- Leave storage space for additional VCDs

  Leave at least 5 GB of free space to allow room to create VCDs for installing additional VMs and applications. (To conserve this storage space, consider deleting VCDs when you are finished using them.)

- Create separate boot and data volumes for each VM

  Install the guest operating system and applications in the first (boot) volume, and create separate volumes for associated data. Separating the boot and data volumes helps to preserve the data and makes it easier to recover a VM if the boot volume crashes.

- Create a boot volume with enough capacity for the guest operating system plus overhead

  Observe the minimum space requirements of your guest operating system and consider allocating slightly more space to account for the formatted capacity of the volume and usage. For example, if you allocate 5 GB to the boot drive when creating the VM, the formatted capacity of the boot volume starts at approximately 4.8 GB before usage, and this might be insufficient to meet a 5 GB requirement.

- Observe the maximum volume size

  When exporting, importing, or restoring a volume, note the maximum volume size, as listed in Important Considerations.

Related Topic

Creating and Migrating Virtual Machines

Managing Virtual Machine Resources

Planning Virtual Machine Networks

Plan network resources to determine how you will allocate available virtual networks to the virtual machines (VMs) on your ztC Edge system.

When you deploy a system licensed for two PMs, software binds pairs of physical network ports across two physical machines (PMs), to form redundant virtual networks. When you create or reprovision VMs on your ztC Edge system, you connect the VMs to virtual networks instead of the physical network ports. For systems licensed for two nodes, software binds pairs of physical network ports across the two physical machines (PMs) to form these redundant virtual networks.

When connecting VMs to virtual networks, consider the following information and restrictions:

- You can connect one VM to multiple virtual networks, and you can connect multiple VMs to the same virtual network.
- The Stratus Redundant Linux software allows unlimited over-provisioning of network resources; therefore, be sure to profile a VM's network bandwidth/response time requirements when allocating virtual networks.
- When multiple VMs share the same virtual network, available network bandwidth is shared equally between the VMs. Unlike vCPU capacity, there is no way to proportionately allocate bandwidth resources. Therefore, high use of network resources by one VM can reduce the performance of all VMs on that network. If a VM has a large bandwidth requirement, consider connecting a dedicated virtual network to that VM.

Related Topics

General Network Requirements and Configurations

Creating and Migrating Virtual Machines

Managing Virtual Machine Resources

Creating and Migrating Virtual Machines

Generate a new virtual machine (VM) on a system by creating a new VM, migrating an existing VM or physical machine (PM) directly over the network, or importing an Open Virtualization Format (OVF) file or Open

Virtualization Appliance (OVA) file from an existing VM.

To create a new VM (without an existing source VM or PM), see Creating a New Virtual Machine.

To copy an existing VM on a system for the purpose of creating a new VM or creating a duplicate VM for troubleshooting, see Copying a Virtual Machine.

To migrate or import a VM from another system, or to restore a VM on the same system, see one of the following topics:

- Migrating a Physical Machine or Virtual Machine to a System

  Use the *P2V client* (**virt-p2v**) to transfer a PM or VM directly over the network to a new VM on the system.

- Exporting a Virtual Machine

  Use the ztC Console to export the source VM to OVF and VHD files on a network share.

- Importing an OVF File

  Use the ztC Console to import OVF and VHD files from another ztC Edge system to the ztC Edge system.

- Replacing/Restoring a Virtual Machine from an OVF File

  Use the ztC Console to import OVF and VHD files back to the same system to overwrite and restore an existing VM from a previous backup copy.

Related Topics

Managing Virtual Machines

Creating a New Virtual Machine

Create a new virtual machine (VM) to install a guest operating system on your ztC Edge system. (You can also migrate an existing VM or physical machine (PM), as summarized in Creating and Migrating Virtual Machines.)

Launch the **VM Creation Wizard** by clicking **Create** on the **Virtual Machines** page. The wizard steps you through the process of allocating resources to the VM.

Prerequisites:

- Review the prerequisites and considerations for allocating CPUs, memory, storage, and network resources to the VM, as listed in Planning Virtual Machine Resources as well as Virtual Machine Recommendations and Limits.
- You can create VMs that run supported guest operating systems, as described in Compatible Guest Operating Systems.
- You can select a remote ISO or a bootable virtual CD (VCD) as the source that the VM boots from. For a remote ISO, you must have a URL or path name for the repository; and for a remote ISO on a shared network drive, you must have a user name and password. If you need a bootable VCD of the Windows or Linux installation media, create it as described in Creating a Virtual CD. The bootable VCD must be a single CD or DVD. Multiple CDs or DVDs are not supported.
- Ensure that both PMs of the ztC Edge system are online and are connected to the network; otherwise, the system cannot properly create the VM.

To create a new VM

1. On the **Physical Machines** page (see The Physical Machines Page) of a system licensed for two nodes, verify that both PMs are in the **running** state and that neither PM is in maintenance mode or in the process of synchronizing.
2. On the **Virtual Machines** page (see The Virtual Machines Page), click **Create** to open the **VM Creation Wizard**.
3. On the **Name, Description, Protection and OS** page:
   a. Type the **Name** and an optional **Description** for the VM as they will appear in the ztC Console.

      The VM name must meet the following requirements:

      - A VM name must start with a word or a number, and the name cannot include the special characters (for example, #, %, or $).
      - A VM name cannot use hyphenated prefixes such as Zombie- or migrating-.
      - A VM name has a maximum of 85 characters.

b. For **Boot From**, select one of the following as the boot source:

- **VCD**—The boot source is a VCD. Select a source from the pull-down menu.
- **Remote ISO via Windows Share (CIFS/SMB)**—The boot source is a remote ISO file on a shared network drive. You must enter values for **User Name** and **Password**. For **Repository**, enter a value in the format ***machine_URL**ShareName*** (for example, **\\192.168.1.34\MyISO_Folder**).
- **Remote ISO via NFS**—The boot source is a remote ISO file, accessed through NFS. For **Repository**, enter the URL of the remote system in the format ***nnn.nnn.nnn.nnn*** (do not include **http://** or **https://**).

For a list of available ISO repositories, click **List ISOs**, and select an ISO file. The full path name of the selected ISO file appears under **Repository**. You cannot edit the ISO URL that is displayed.

c. Click **Next**.

4. On the **vCPUs and Memory** page:
   a. Specify the number of **vCPUs** and the amount of **Memory** to assign to the VM. For more information, see Planning Virtual Machine vCPUs and Planning Virtual Machine Memory.
   b. Click **Next**.
5. On the **Volumes** page:
   a. Type the **Name** of the boot volume as it will appear in the ztC Console.
   b. Type the **Volume Size** of the volume to create in gigabytes (GB). For more information about allocating storage, see Planning Virtual Machine Storage.
   c. If applicable, create additional data volumes by clicking **Add New Volume** and specifying the parameters for each volume. (You can also add volumes after you create the VM by using the **Reprovision Virtual Machine** wizard, as described in Creating a Volume in a Virtual Machine.)
   d. Click **Next**.
6. On the **Networks** page, select the shared networks to attach to the VM (for more information, see Planning Virtual Machine Networks). You can also enable (or disable) the network and specify the MAC address. To continue, click **Next**.
7. On the **Creation Summary** page:

a. Review the creation summary. If you need to make changes, click **Back**.

b. If you want to prevent a console session from automatically starting to observe the software installation, deselect **Launch Console**.

c. To accept the VM as provisioned and begin the software installation, click **Finish**.

The **VM Creation Wizard** displays progress of the creation and opens the console window, if applicable.

8. If applicable, observe the progress of the installation of the operating system (allow pop-ups in your browser, if necessary) and respond to any prompts in the VM console session.
9. After you install the operating system, configure the additional resources and software necessary for production use, as described in:
   - Configuring Windows-based Virtual Machines
   - Configuring Linux-based Virtual Machines

Caution: If the primary PM fails or the VM crashes before the final reboot after the installation process is completed, the installation of the VM may need to be restarted.

The VM may not reboot if installations of any of the following are aborted:

- The guest operating system, including the configuration steps
- Any middleware or applications that manipulate system files

Related Topics

Copying a Virtual Machine

Renaming a Virtual Machine

Removing a Virtual Machine

Creating and Migrating Virtual Machines

Managing Virtual Machine Resources

Managing the Operation of a Virtual Machine

Copying a Virtual Machine

Copy a virtual machine (VM) if you want to clone an existing VM on your ztC Edge system. For example, you can copy a healthy VM to create a new VM, or you can copy a VM that is not working properly and use the

copy for troubleshooting purposes. (If you want to import or migrate a VM from a different system, see the overview in Creating and Migrating Virtual Machines.)

To copy a VM, select a VM on the **Virtual Machines** page and click **Copy**. A wizard steps you through the process of renaming and allocating resources to the new VM.

Copying a VM creates an identical VM with a unique SMBIOS UUID, system serial number, MAC addresses, and hardware ID.

Notes:

- To prevent conflicts with the source VM, the copy wizard automatically assigns a new MAC address to each network interface in the new VM; however, you may need to manually update any IP addresses and host names.
- If the ztC Edge system switches from the primary PM to the secondary PM while copying a VM, the copy process fails. This does not affect the continuous uptime of your system, but you must delete any volumes associated with the copied VM and start the copy again.
- For information on the modes of operation that systems support, see Virtual Machine Recommendations and Limits.
- You can copy a VM on a system licensed for one node.

Prerequisites:

- You must shut down a VM before copying it.
- On a system licensed for two nodes, both PMs of the ztC Edge system must be online for the copy process to function properly.

To copy a VM on the ztC Edge system

1. On the **Physical Machines** page (see The Physical Machines Page) of a system licensed for two nodes, verify that both PMs are in the **running** state and that neither PM is in maintenance mode or in the process of synchronizing. On a system licensed for one node, verify that the PM is in the **running** state.
2. On the **Virtual Machines** page (see The Virtual Machines Page), select the VM that you want to copy and click **Shutdown**.
3. When the VM has stopped, click **Copy** to open the copy wizard.

4. On the **Name, Description, and Protection** page:
    a. Type the **Name** and an optional **Description** for the VM as they will appear in the ztC Console.
    b. Click **Next**.
5. On the **vCPUs and Memory** page:
    a. Specify the number of **vCPUs** and the amount of **Memory** to assign to the VM. For more information, see Planning Virtual Machine vCPUs and Planning Virtual Machine Memory.
    b. Click **Next**.
6. On the **Volumes** page, you can:
    - Type the volume **Name**.
    - Specify the **Volume Size** of each volume.
    - Click **Add New Volume** to create a new data volume. (If the button is not visible, scroll down to the bottom of the wizard page.)

    For more information, see Planning Virtual Machine Storage. To continue, click **Next**.
7. On the **Networks** page, activate the check box for each shared network that you want to attach to the VM.
8. On the **Copy Summary** page:
    a. Review the configuration summary. If you need to make changes, click **Back**.
    b. To proceed with copying the VM, click **Finish**.

After the copy process is complete; the ztC Edge system may continue to synchronize data between PMs to enable HA operation.

Troubleshooting

If necessary, use the following information to resolve problems with the copy process.

To clean up after a canceled or failed copy process

Remove any volumes associated with the copied VM.

Related Topics

Creating and Migrating Virtual Machines

Managing Virtual Machine Resources

Managing the Operation of a Virtual Machine

Migrating a Physical Machine or Virtual Machine to a System

Migrate a physical machine (PM) or virtual machine (VM) to transfer it over an A-link network to a new VM on the system. (You can also import an Open Virtualization Format (OVF) file to a system, as summarized in Creating and Migrating Virtual Machines.)

Procedures below describe how to migrate a PM or VM over the network: download the *P2V client* (**virt-p2v**) ISO file, boot the P2V client ISO file on the source PM or VM, and then use the client to configure, initiate, and monitor the secure network transfer from source side. No configuration steps are required on the system until after the migration is complete, but you can confirm that the migration is in progress on the **Volumes** page of the ztC Console as volumes associated with the new VM begin to appear.

Caution: Consider backing up the source PM or VM before preparing to migrate it. To backup a VM, export it (see Exporting a Virtual Machine).

Notes:

- The migration process supports PMs or VMs running only the following operating systems:
  - CentOS/RHEL 7.4 or 7.5
  - Microsoft Windows 7, 8.x, or 10; or Windows Server 2008 R2, 2012, or 2016.
  - Ubuntu 14.04 LTS, 16.04 LTS, or 18.04 Server–After migrating this VM, you need to perform additional procedures. See To complete the migration of an Ubuntu VM.
  - VMware Release 6.x
- For Windows-based VMs that support *hibernation* or *fast startup* mode, you must disable these features before the migration process. To fully disable hibernation or fast startup mode, see the instructions to recover from a migration that fails with the error message `Failed to mount '/dev/sda1: Operation not permitted` below in **Troubleshooting**.
- For Linux-based PMs or VMs, consider editing the `/etc/fstab` file before the migration process to comment out entries for data volumes and allow only the boot volume to mount. Because Linux-based VMs use different device names on the ztC Edge system, a new VM may boot into single-user mode if it cannot mount volumes with their original device names. You can restore the `/etc/fstab` entries with the correct device names after the migration, as described below in **Troubleshooting**.
- When migrating a VMware VM, you must shutdown the VM using operating system shutdown commands in addition to powering it off from the VMware console. If you shutdown the VM using only the VMware console, the migration will fail.
- The source PM or VM must be offline for the duration of the migration process. Consider scheduling a planned maintenance period for the migration.
- The time required for the migration depends on the size and number of volumes on the source system as well as the network bandwidth between the source and the target system. For example, transferring a source system with one 20 GB boot volume over a 1Gb network may take about 30 minutes.
- You can migrate multiple PMs or VMs at one time, but sharing network bandwidth may increase migration times.

- To prevent conflicts with the original PM or VM, the P2V client automatically assigns a new MAC address to each network interface in the new VM; however, you must manually update any IP addresses and host names as needed.
- If the system switches from the primary PM to the secondary PM during a migration, the migration process fails. This does not affect the continuous uptime of your system, but you must reboot the P2V client on the source PM or VM and start over. See **Troubleshooting** below for more information.
- After migrating a PM or VM, the network driver might not be properly installed. In this situation, manually install the driver. See **Troubleshooting** below for more information.

Prerequisite: For a system licensed for two nodes, both PMs of the system must be online for the migration process to function properly (you cannot migrate PMs or VMs on a system licensed for one node). On the **Physical Machines** page of the ztC Console, verify that both PMs are in the **running** state and that neither PM is in maintenance mode or in the process of synchronizing.

Perform the following migration procedures (click drop-down menus, if applicable).

To prepare for migrating a PM to the ztC Edge system

1. Download the P2V client ISO file from the **Downloads** page at https://www.stratus.-com/services-support/downloads/?tab=ztcedge.
   a. On the **Downloads** page, click **ztC Edge** (if it is not already displayed) and then select the appropriate version.
   b. Scroll down to **Drivers and Tools** and then continue scrolling to **ztC Edge P2V Client for Virtual or Physical Machine Migration**.
   c. Select the **P2V Client (virt-p2v)** file.
2. If you want to verify the integrity of the ISO image, use the associated fciv checksum file and the Microsoft File Checksum Integrity Verifier (FCIV) executable file, if you have already downloaded the Microsoft file to your system. Otherwise, use the MD5 checksum hash function.

   Using fciv

   Download the associated fciv verification file from the **Drivers and Tools** section of the **Downloads** page. Save the fciv verification file to the directory that contains the ISO file you downloaded.

Open a command prompt window. From the directory containing the ISO, executable, and verification files, type a command similar to the following to verify the ISO image:

fciv -v -xml virt-p2v-*n.n.n-n*.nnnnnnnn.n.el6.centos.xml

If the command succeeds (that is, it returns the message All files verified successfully), go to the next step. If the command fails, repeat the download.

Using MD5

Open a command prompt window as an administrator, and enter the following:

CertUtil -hashfile *path_to_file* MD5

The **CertUtil** command displays a message indicating whether or not it completed successfully. If the command succeeds, continue with the next step. If the command fails, repeat the download.

3. Burn the P2V client ISO file to a CD-ROM that you will use to boot the source PM.
4. Insert the P2V client CD into the CD/DVD drive of the source PM.
5. Shut down the PM in preparation to boot the P2V client.

To prepare for migrating a VM to the ztC Edge system

1. Download the P2V client ISO file from the **Drivers and Tools** section of the **Downloads** page at https://www.stratus.com/services-support/downloads/?tab=ztcedge. Ensure that you download the version of the P2V client that matches the version of the ztC Edge system to which you are migrating the VM.
2. If you want to verify the integrity of the ISO image, use the associated fciv checksum file and the Microsoft File Checksum Integrity Verifier (FCIV) executable file, if you have already downloaded the Microsoft file to your system. Otherwise, use the MD5 checksum hash function.

   Using fciv

   Download the associated fciv verification file from the **Drivers and Tools** section of the **Downloads** page. Save the fciv verification file to the directory that contains the ISO file you downloaded.

   Open a command prompt window. From the directory containing the ISO, executable, and verification files, type a command similar to the following to verify the ISO image:

   fciv -v -xml virt-p2v-*n.n.n-n*.nnnnnnnn.n.el6.centos.xml

If the command succeeds (that is, it returns the message All files verified successfully), go to the next step. If the command fails, repeat the download.

Using MD5

Open a command prompt window as an administrator, and enter the following:

CertUtil -hashfile *path_to_file* MD5

The **CertUtil** command displays a message indicating whether or not it completed successfully. If the command succeeds, continue with the next step. If the command fails, repeat the download.

3. Insert or connect the P2V client ISO file to the source VM and set the virtual CD drive as the boot device in the associated hypervisor.
4. Shut down the VM in preparation to boot the P2V client.

To migrate a PM or VM to the ztC Edge system

1. Power on the source PM or VM to boot the P2V client. After a minute or so, the **virt-p2v** window is displayed.
2. The P2V client automatically obtains network settings through DHCP. Static settings are unnecessary for the migration process, but you can optionally click **Configure network** to specify the settings. (If necessary, configure the network settings of the target VM later on the ztC Edge system.)
3. Enter the connection settings for the **Conversion server** (the ztC Edge system). Enter the hostname or IP address of the system and the **Password** for the `root` account. (You must use the `root` account of the ztC Edge host operating system, as described in Accessing the Host Operating System.)
4. Click **Test connection**. If the P2V client connects to the ztC Edge system, click **Next** to continue. A page appears with sections for **Target properties**, **Fixed hard disks**, and other settings.

   If the P2V client cannot connect, verify the connection settings and try to connect again.
5. In the **Target properties** section, enter the **Name** for the target VM that will be displayed in the ztC Console. (The name must be different from any existing VMs on the ztC Edge system.)

6. The **# vCPUs** and **Memory(MB)** values are automatically detected and completed, but optionally modify them if you want the VM on the ztC Edge system to have more CPUs or memory than the source PM or VM.
7. Specify the **Virt-v2v output options** for the target VM, as follows:
    a. Next to **Output to**, select **everrunha** (High Availability) operation. (For information about operation options, see Creating a New Virtual Machine and Modes of Operation.)
    b. Next to **Output format**, select the disk image format, **raw** or **qcow2**.
8. If you want to save debugging messages from the migration process, optionally select the **Enable server-side debugging** check box. (The debugging messages are included if you generate a diagnostic file for your authorized Stratus service representative, as described in Creating a Diagnostic File.)
9. Select which **Fixed hard disks** (volumes) to include in the migration by activating the check box next to each device.

    You must select at least one volume, including the boot volume. (Because the P2V client is a Linux-based utility, all devices are listed by Linux device names, where **sda** or **vda** represents the boot volume.)
10. Select which **Network Interfaces** to include in the migration by activating the check box next to each device.

    If the target ztC Edge system has more than one shared network, you can also select the shared network to connect with each network interface. Double-click the network interface to open the **Configure Network** dialog box and select the shared network from a drop-down list.

    In the **Configure Network** dialog box, you can also specify a MAC address for a specific network interface. If you do not specify an address, the system automatically sets the MAC address for each network interface.

    Click **OK** when you have finished configuring the network interface.
11. When you are ready to migrate the PM or VM to the ztC Edge system, click **Start conversion**. (If you need to cancel the migration for any reason, see **Troubleshooting** below.)
12. When the migration is complete, the P2V client displays a success message. If applicable, you can eject the CD or virtual CD and click **Power Off** to shut down the source PM or VM.

Note: After the migration, the new VM on the ztC Edge system is located on the primary PM, and it remains in a stopped state. Before starting the VM, complete the migration as described in the next procedure.

To complete the migration on the ztC Edge system

1. Open the **Virtual Machines** page (see The Virtual Machines Page) in the ztC Console.
2. Select the new VM in the top pane and click **Config** to open the **Reprovision Virtual Machine** wizard, as described in Reprovisioning Virtual Machine Resources. Use the wizard to configure the desired vCPUs, memory, storage, and network settings for the VM:
   - If your source PM or VM had more than one network interface, configure the additional network interfaces that were not included in the migration process.
   - If you will continue running the source PM or VM, ensure that the MAC address for each network interface in the new VM is different from the source PM or VM.

   Click **Finish** on the last wizard page to implement the changes.
3. Click **Start** to boot the new VM.
4. Click **Console** to open the console of the VM and log on to the guest operating system. (For information about using the console, see Opening a Virtual Machine Console Session.)
5. Disable any guest operating system services that are unnecessary for operation on the ztC Edge system:
   - If you migrated from a PM source, disable any services that interact directly with hardware. Examples include:
     - Dell OpenManage (OMSA)
     - HP Insight Manager
     - Diskeeper
   - If you migrated from a VM source, disable any services associated with other hypervisors. Examples include:
     - VMware Tools
     - Hyper-V Tools
     - Citrix Tools for Virtual Machines

After disabling these services, restart the guest operating system to implement your changes.

6. If necessary, update the network configuration settings in the guest operating system and restart it to enable the settings.
7. Verify that you have configured your guest operating system with the additional Windows- or Linux-based system settings described in:
   - Configuring Windows-based Virtual Machines
   - Configuring Linux-based Virtual Machines

After you verify that the new VM is functioning properly, the migration process is complete; however, the system may continue to synchronize data between PMs to enable High Availability (HA) operation.

To complete the migration of an Ubuntu VM

After migrating a VM using P2V from a bare metal machine running an Ubuntu release, the VM may have problems, such as no active network. To correct the problem, perform the appropriate procedure below after migrating the Ubuntu VM.

After migrating an Ubuntu 16.04 VM

1. From the ztC Console, open a console window into the VM.
2. Log in to the VM and go to the terminal.
3. Enter the following command: `cd /etc/network`.
4. Enter the following command: `sudo vi interfaces`.
5. Change the `eno1` interface to `ens3f0`.
6. Save the changes.
7. Enter the following command: `sudo systemctl restart networking.service`.

After migrating an Ubuntu 14.04 VM

The following procedure does not require you to reboot the system.

1. From the ztC Console, open a console window into the VM.
2. Log in to the VM and go to the terminal.
3. Enter the following command: `ifconfig`.
   Note that the command output does not include `eth0`.

4. Enter the following command:. `ifconfig -a`.

   Note that the command output includes `eth0`.

5. Enter the following command: `cd /etc/network/`.
6. Enter the following command: `sudo vi interfaces`.

   In the `interfaces` file, change `em1` to `eth0`.

7. Enter the following command: `cd run`.
8. Enter the following command: `ls -l`.
9. Enter the following command: `sudo touch ifup.eth0`.
10. Enter the following command: `sudo vi ifstate.eth0`.

    In the `ifstate.eth0` file, insert `eth0` at the top of the file.

11. Enter the following command: `sudo ifdown eth0`.
12. Enter the following command: `sudo ifup eth0`.

The following procedure requires you to reboot the VM:

1. From ztC Console, open a console window into the VM.
2. Log in to the VM and go to the terminal.
3. Edit the `/etc/network/interfaces` file, changing `em1` to `eth0`.
4. Reboot the VM.

Troubleshooting

If necessary, use the following information to resolve problems with the migration process.

To cancel the migration process

Power down the source PM or VM running the P2V client.

To clean up after a canceled or failed migration

Open the ztC Console and remove any migrated volumes associated with the source PM or VM. If you want to restart the migration process, reboot the P2V client on the source PM or VM.

To recover from a failed migration

If the migration process fails, an error message is displayed in the P2V client on the source PM or VM. Another message may be displayed on the ztC Edge system. Use these messages to determine the problem.

If the migration continues to fail, and the option is available, enable server-side debugging. After the migration, generate a diagnostic file to send to your authorized Stratus service representative, as described in Creating a Diagnostic File. The diagnostic file includes any server-side debugging messages from the migration process.

To recover from a migration that fails with the error message, `Failed to mount '/dev/sda1: Operation not permitted`

For Windows-based PMs or VMs, if the migration process fails with the following error message, it may indicate that *hibernation* or *fast startup* mode are enabled:

```
Failed to mount '/dev/sda1': Operation not permitted
The NTFS partition is in an unsafe state. Please resume and
shutdown Windows fully (no hibernation or fast restarting), or
mount the volume read-only with the 'ro' mount option.
```

To resolve the issue, disable hibernation and fast startup in the source PM or VM:

1. Log on to the operating system of the source PM or VM.
2. Open the **Power Options** control panel and click **Choose what the power buttons do**.
3. Next to **When I press the power button**, select **Shutdown** (instead of **Hibernate** or **Sleep**, if present).
4. Under **Shutdown Settings**, clear the check box next to **Turn on fast startup (recommended)**, if present.
5. Click **Save changes**.
6. Open **Administrator Power Shell** and execute the following command:

   ```
   > powercfg /h off
   ```

7. Shut down the operating system and restart the migration process.

To recover when a newly migrated Linux-based VM is stuck in the "booting" state

A Linux-based VM may fail to exit the **booting** state in ztC Console if the VM's network is offline.

During the migration process, the P2V client attempts to set a new MAC address for each network interface to prevent conflicts with the original VM. Some Linux-based operating systems detect a new MAC address and automatically create a new network interface for it while still retaining the original interface. The guest operating system boots, but the network may remain offline until you manually configure the network settings.

To correct the problem, open the VM console, log on to the guest operating system, and update the network startup scripts. Ensure that you retain only one entry for each network interface, and that each interface uses a unique MAC address and correct network settings for your environment.

To recover missing data volumes in the VM on the ztC Edge system

If the data volumes do not appear in the VM on the ztC Edge system after the import, you may need to manually restore the volumes, as follows:

- Shut down the VM, run the **Reprovision Virtual Machine** wizard, and verify that you have included the volumes on the **Volumes** page.
- For Windows-based VMs, use **Disk Management** to bring data volumes online.
- For Linux-based VMs, edit the `/etc/fstab` file to reflect the new device names for the storage devices (`/dev/vda` through `/dev/vdh`). Device names also may have shifted, for example, if volumes were not included in the import.

To recover missing network devices in the VM on the ztC Edge system

If the network devices do not appear in the VM on the ztC Edge system after the import, you may need to manually restore them, as follows:

- Shut down the VM, run the **Reprovision Virtual Machine** wizard, and verify that you have included the networks on the **Networks** page.
- For Linux-based VMs, reconfigure the network startup script to reflect the new device names for the network interfaces.

To manually install a new network driver

After migrating a PM or VM, the network driver might not be properly installed (for example, Device Manager might list the driver with a warning, ). In this situation, manually install the driver:

1. In the VM console window, open **Device Manager** in the guest operating system.
2. Expand **Network adapters** and right-click the **Red Hat VirtIO Ethernet Adapter** (the driver that does not work correctly).
3. Select **Update Driver Software**.
4. In the pop-up window, click **Browse my computer for the driver software**.
5. Click **Let me pick from a list of device drivers**.
6. Select **Red Hat VirtIO Ethernet Adapter**.
7. Click **Next** to install the network driver.

After the driver is installed, check the VM's state in the ztC Console. If the state is running (), the driver is working properly.

Related Topics

Creating and Migrating Virtual Machines

Configuring Windows-based Virtual Machines

Configuring Linux-based Virtual Machines

Managing Virtual Machine Resources

Managing the Operation of a Virtual Machine

Importing an OVF File

Import an Open Virtualization Format (OVF) file from a system if you want to transfer a VM from one system another, or if you want to transfer an image that you created back to the same system to restore or duplicate the original VM. (To migrate a physical machine (PM) or virtual machine (VM) to a system without using an OVF file, see Migrating a Physical Machine or Virtual Machine to a System.)

You can *import* or *restore* the VM. Importing a VM creates a new instance of the VM with unique hardware IDs. Restoring a VM creates an identical VM with the same hardware IDs (SMBIOS UUID, system serial number, and MAC addresses, if provided in the VM image) that your guest operating system and applications may require for software licensing. To prevent conflicts with the original VM, restore a VM only if you want to transfer it to the ztC Edge system and stop using it on the source system.

This topic explains how to import an OVF file from a local computer, a USB device, or a remote file system such as an NFS export or a Windows share (also known as a CIFS share, such as, for example, Samba).. If

you want to restore an existing VM on the same system to overwrite the VM and recover it from a previous backup copy, see Replacing/Restoring a Virtual Machine from an OVF File.

Notes:

- Import a VM if you are trying to create or clone a VM from a golden image, since the system will assign unique hardware ID and MAC addresses when importing a VM. (A golden image is typically a template VM created for the purpose of copying multiple times.) To prevent conflicts with the source VM, the import wizard automatically assigns a new MAC address to each network interface in the new VM; however, you may need to manually update any IP addresses and host names as needed.
- You can import VMs only if they are running supported guest operating systems, as described in Compatible Guest Operating Systems.
- When you import a VM back to the same system to duplicate the VM, you must rename the VM and duplicate volumes during either the export or import process. If you do not rename the VM, the import wizard automatically renames the new VM and new volumes, to prevent conflicts with the source VM. The wizard appends a number to the VM name and volume name, incrementing the number for additional duplicates of the VM: **MyVM**, **MyVM0**, **MyVM1**, and so on.
- The time required to import a VM depends on the size and number of volumes in the source VM as well as network bandwidth. For example, transferring a VM with one 20 GB boot volume over a 1Gb network may take about 30 minutes.
- If the system switches from the primary PM to the secondary PM during an import process, the process fails. This does not affect the continuous uptime of your system, but you must delete the incomplete VM and associated volumes on the system, and import them again.
- After migrating a PM or VM, the network driver might not be properly installed. In this situation, manually install the driver. See **Troubleshooting** below for more information.

Prerequisite:
Before you import a VM image from an OVF file, use the ztC Console on the source system to export a VM (see Exporting a Virtual Machine) to OVF and Virtual Hard Disk (VHD) files on a supported network share or a USB device. Copy these files to your management PC, or mount the USB device or network share on the target ztC Edge system as described in Mounting a USB Device or Network-mounted Folder on the ztC Edge System, and then use the ztC Console on the target system to import the OVF and VHD files.

To import an OVF file

1. Log on to the ztC Console on the target system.
2. On the **Physical Machines** page (see The Physical Machines Page) of a system licensed for two nodes, verify that both PMs are in the **running** state and that neither PM is in maintenance mode or in the process of synchronizing.
3. If you are importing a VM from a USB device or network share (instead of the PC running the ztC Console), mount the device or share on the ztC Edge system as described in Mounting a USB Device or Network-mounted Folder on the ztC Edge System.
4. On the **Virtual Machines** page (see The Virtual Machines Page), click **Import/Restore** to open the **Import/Restore Virtual Machine** wizard.
5. Select one of the following:
   - **Browse for OVF**

     Note: Browsing for VMware OVFs is not supported.

     Click **Browse** and locate the OVF on a local computer.
   - **Browse for OVF from USB**

     Select a partition from the pull-down menu, click **List OVFs**, and select a file from the pull-down menu.
   - **Remote OVF via Windows Share(CIFS/SMB)**

     The OVF is located in a folder on a CIFS share. Enter values for **Username** and **Password**. For **Repository**, enter a value in the format *machine_URL**ShareName* (for example,

\\192.168.1.34\MyOVFsForImport). Then, click **List OVFs** and select an OVF file from the list.

- **Remote OVF via NFS**

  The OVF is in a folder on a remote system accessed through NFS. For **Repository**, enter the URL of the remote system in the format ***nnn.nnn.nnn.nnn*** (do not include **http://** or **https://**). Click **List OVFs** to display a list of all OVF files in the remote folder. Select the file to import.

6. Select **Import** (scroll down the window, if necessary). Import creates a new instance of the VM with unique hardware IDs.
7. When prompted, click **Browse** and the wizard displays the **Importing Virtual Machine** window. Select the appropriate **.vhd** file to include for each volume associated with the VM.
8. Review and, if necessary, edit the information (you may need to scroll down the window):

   - **Name, CPU, and Memory**

     Displays the name of the VM, the number of vCPUs, and the total memory the VM can use. Edit the information, if necessary.

   - **Storage**

     Displays the name and size of each volume. In the **Create** column, select a box for a volume to allocate storage for the volume on the system (the boot volume is required). In the **Restore Data** column, select a box to import data for a volume from the VHD file.

   - **Network**

     Displays the available networks. You can remove a network or add one that is not already allocated. You can also specify a MAC address for each selected network. A minimum of one network is required.

9. Optionally, clear the check box for **Auto start Virtual Machine after import** if you need to reprovision the VM before starting it for the first time.
10. Click **Import** to begin importing the VM.

    The wizard displays progress information. When the transfer is complete, click **Done** to close the wizard.

> **Note**: Imported volumes begin to appear on the **Volumes** page of the ztC Console while the import is still in progress. Do not attach or remove any of these imported volumes until the import window reports that the process is complete; otherwise, the import fails.

11. If applicable, use the **Reprovision Virtual Machine** wizard to allocate additional resources to the VM, as described in Reprovisioning Virtual Machine Resources.

    When you are finished reprovisioning the VM, click **Start** to boot the VM.

12. Click **Console** to open the console of the VM and log on to the guest operating system.
13. For Windows-based VMs only, download and update the VirtIO drivers to the latest supported versions, as described in Updating the VirtIO Drivers (Windows-based VMs). (The correct VirtIO drivers are already present in Linux-based VMs.)

    > **Note**: After updating the drivers, you may need to restart the guest operating system.

14. If necessary, update the network settings in the guest operating system.

After you verify that the new VM is functioning properly, the import process is complete; however, the system may continue to synchronize data between PMs to enable High Availability (HA) operation.

> **Note**: The new VM and its associated volumes may be marked with warning symbols until the data has been synchronized and the VirtIO drivers are running.

Troubleshooting

If necessary, use the following information to resolve problems with the export or import process.

To clean up after a canceled or failed import

In the ztC Console on the target system, remove the imported VM and any volumes associated with the imported VM, if present.

To recover missing data volumes in the target VM

If data volumes do not appear in the VM on the target system after the import, you may need to manually restore the volumes, as follows:

- Shut down the VM, run the **Reprovision Virtual Machine** wizard, and verify that you have included the volumes on the **Volumes** page.

- For Windows-based VMs, use **Disk Management** to bring data volumes online.
- For Linux-based VMs, edit the `/etc/fstab` file to reflect the new device names for the storage devices. Device names may have shifted, for example, if volumes were not included in the import.

To recover missing network devices in the VM on the ztC Edge system

If network devices do not appear in the VM on the target system after the import, you may need to manually restore them, as follows:

- Shut down the VM, run the **Reprovision Virtual Machine** wizard, and verify that you have included the networks on the **Networks** page. If the VM requires more networks than shown in the wizard, connect additional business networks to the ztC Edge system and then reprovision the VM to include the new networks.
- For Linux-based VMs, reconfigure the network startup script to reflect the new device names for the network interfaces.

To manually install a new network driver

After importing a PM or VM, the network driver might not be properly installed (for example, Device Manager might list the driver with a warning, ). In this situation, manually install the driver:

1. In the VM console window, open **Device Manager** in the guest operating system.
2. Expand **Network adapters** and right-click the **Red Hat VirtIO Ethernet Adapter** (the driver that does not work correctly).
3. Select **Update Driver Software**.
4. In the pop-up window, click **Browse my computer for the driver software**.
5. Click **Let me pick from a list of device drivers**.
6. Select **Red Hat VirtIO Ethernet Adapter**.
7. Click **Next** to install the network driver.

After the driver is installed, check the VM's state in the ztC Console. If the state is running (), the driver is working properly.

Related Topics

Mounting a USB Device or Network-mounted Folder on the ztC Edge System

Creating and Migrating Virtual Machines

Configuring Windows-based Virtual Machines

Configuring Linux-based Virtual Machines

Managing Virtual Machine Resources

Managing the Operation of a Virtual Machine

Replacing/Restoring a Virtual Machine from an OVF File

Replace a virtual machine (VM) from a ztC Edge-created Open Virtualization Format (OVF) file if you want to restore (that is, recover) a VM on your ztC Edge system by overwriting the VM with a previous backup copy. (If you want to import a VM from a different system, see the overview in Creating and Migrating Virtual Machines.)

Typically, importing a VM creates a new instance of the VM with unique hardware IDs. Restoring a VM creates an identical VM with the same SMBIOS UUID, system serial number, and MAC addresses, if provided in the VM image, that your guest operating system and applications may require for software licensing. The hardware ID, though, of the restored VM is unique. If an identical VM already exists on the ztC Edge system, restoring the VM allows you to replace the VM and overwrite it with your previous copy.

You can restore a VM that already exists on a ztC Edge system only if you have previously exported a VM (see Exporting a Virtual Machine) from a ztC Edge system to OVF and Virtual Hard Disk (VHD) files on a supported network share or a USB device. Copy these files to your management PC, or mount the USB device or network share on the target ztC Edge system as described in Mounting a USB Device or Network-mounted Folder on the ztC Edge System, and then use the ztC Console on the target ztC Edge system to restore the OVF and VHD files from your management PC.

Caution: Consider backing up your existing VM on the ztC Edge system before overwriting and restoring it. If you export the VM to create the backup, ensure that you do not overwrite the OVF and VHD files that you want to restore.

Notes:

- You can restore a VM from only an OVF created from a ztC Edge system. You cannot restore a VM from an OVF created from a third-party system.
- You typically restore a VM to recover the VM from a previous backup. When restoring a VM, the system attempts to preserve the hardware ID and MAC addresses of all network interfaces.
- Restore a VM only if you are specifically trying to restore a particular instance of a ztC Edge VM and that the restored VM will be the only copy of this VM running across all ztC Edge servers in your network.
- The time required to restore a VM depends on the size and number of volumes in the source VM as well as network bandwidth. For example, transferring a VM with one 20 GB boot volume over a 1Gb network may take about 30 minutes.
- If you overwrite and restore an existing VM, the ztC Edge system removes the existing VM and its volumes.
- If the ztC Edge system switches from the primary PM to the secondary PM while restoring a VM, the restore process fails. This does not affect the continuous uptime of your system, but you must delete the incomplete VM and associated volumes on the ztC Edge system, and restore them again.

Prerequisites:

- Before you replace (that is, restore) a VM image from a ztC Edge system, use the ztC Console on the source ztC Edge system to export a VM (see Exporting a Virtual Machine) to OVF and Virtual Hard Disk (VHD) files on a supported network share or a USB device. Copy these files to your management PC, or mount the USB device or network share on the target ztC Edge system as described in Mounting a USB Device or Network-mounted Folder on the ztC Edge System, and then use the ztC Console on the target ztC Edge system to restore the OVF and VHD files
- Both PMs of the ztC Edge system must be online for the restore process to function properly.

To restore a VM

1. Log on to the ztC Console on the target ztC Edge system.
2. On the **Physical Machines** page (see The Physical Machines Page) of a system licensed for two nodes, verify that both PMs are in the **running** state and that neither PM is in maintenance mode or in the process of synchronizing.
3. If you are restoring a VM from a USB device or network share (instead of the PC running the ztC Console), mount the device or share on the ztC Edge system as described in Mounting a USB Device or Network-mounted Folder on the ztC Edge System.
4. On the **Virtual Machines** page (see The Virtual Machines Page), select the VM that you want to restore in the upper panel.
5. In the lower panel, click **Restore** or click **Import/Restore** near the top pane.
6. Select one of the following:
   - **Browse for OVF**

     Click **Browse** and locate the OVF on a local computer.
   - **Browse for OVF from USB**

     Select a partition from the pull-down menu, click **List OVFs**, and select a file from the pull-down menu.
   - **Remote OVF via Windows Share(CIFS/SMB)**

     The restore source is a folder on a CIFS share. Enter a **Username**, **Password**, and **Repository** value. For **Repository**, enter a value in the format ***machine_URL**ShareName*** (for example, **\\192.168.1.34\MyOVFs**). Then, click **List OVFs** and select an OVF file from the list.
   - **Remote OVF via NFS**

     The restore source is a folder on a remote system accessed through NFS. For **Repository**, enter the URL of the remote system in the format ***nnn.nnn.nnn.nnn*** (do not include **http://** or **https://**). Then, click **List OVFs** to display a list of all OVF files in the remote folder. Select the file to import or restore.
7. Select **Restore**. (Scroll down the window, if necessary.) A warning message appears.
8. Click **Continue** to proceed.

> **Caution**: Restoring a VM overwrites all of its data and configuration details.

9. If prompted, add VHD files.
10. Review the information and make any desired edits, if necessary:
    - **Name, CPU, and Memory**

      Displays the name of the VM, the number of vCPUs, and the total memory the VM can use. Edit the information, if necessary.

    - **Storage**

      Displays the name and size of each volume. In the **Create** column, select a box for a volume to allocate storage for the volume on the ztC Edge system (the boot volume is required). In the **Restore Data** column, select a box to import data for a volume from the VHD file.

    - **Network**

      Displays all of the available networks. You can remove a network or add one that is not already allocated. A minimum of one network is required.
11. Optionally, clear the check box for **Auto start Virtual Machine after restore** if you need to reprovision the VM before starting it for the first time.
12. Click **Restore** to begin restoring the VM. When the transfer is complete, click **Done** to close the wizard.

    > **Note**: Restored volumes begin to appear on the **Volumes** page of the ztC Console while the restore process is still in progress. Do not attach or remove any of these restored volumes until the restore window reports that the process is complete; otherwise, the restore process fails.
13. If applicable, use the **Reprovision Virtual Machine** wizard to allocate additional resources to the VM, as described in Reprovisioning Virtual Machine Resources.

    When you are finished reprovisioning the VM, click **Start** to boot the VM.

After you verify that the restored VM is functioning properly, the restore process is complete; however, the ztC Edge system may continue to synchronize data between PMs to enable High Availability (HA) operation.

> **Note**: Your restored VM and its associated volumes may be marked with warning symbols until the data has been synchronized and the VirtIO drivers are running.

Troubleshooting

If necessary, use the following information to resolve problems with the restore process.

To clean up after a canceled or failed restore process

In the ztC Console on the target system, remove the restored VM and any volumes associated with the restored VM, if present.

Related Topics

Creating and Migrating Virtual Machines

Managing Virtual Machine Resources

Managing the Operation of a Virtual Machine

Exporting a Virtual Machine

Export a virtual machine (VM) from a system in order to save an image of the VM to a network-mounted folder (that is, directory) or to a USB device. Exporting a VM from an ztC Edge system makes the VM image available for importing to another system or for importing back to the same ztC Edge system to restore or duplicate the original VM. An exported VM can function as a backup of the original VM. You can directly export a VM from the ztC Edge system as described in this topic.

> **Note**: Exporting a VM as a backup is particularly important for VMs on a system licensed for one node. In the event of a failure or loss of the original VM, you can use the exported VM to restore the lost VM.

Prepare for exporting a VM by inserting a USB device or by creating a network-mounted folder to store an exported VM in your environment. If you are using a USB device, insert it into the current primary node of the system (displayed as **node*n* (primary)** on the **Physical Machines** page). If you are using a folder, create a folder for either a Windows share or a Network File System (NFS) export. A Windows share is also known as a Common Internet File System (CIFS) share (Samba, for example). Then mount the folder or USB device in the host operating system of the ztC Edge system, as described in this topic. When you initiate an export in the ztC Console, the ztC Edge system saves the VM as standard Open Virtualization Format (OVF) and Virtual Hard Disk (VHD) files.

Notes:

- Because the source VM must be shut down to export is, consider scheduling a planned maintenance period for this process.
- The time required for the export depends on the size and number of volumes in the source VM as well as network bandwidth. For example, transferring a VM with one 20 GB boot disk over a 1Gb network may take about 30 minutes.
- If you will continue to use the source VM after exporting it, remember to set a different MAC address and IP address for the VM when you import it on the target system.
- If the ztC Edge system switches from the primary PM to the secondary PM during an export, the process fails. This does not affect the continuous uptime of the system. You can delete the partially exported files from the network-mounted folder and export the files again.
- The maximum size of a vfat file that you can export is 4 GB. If you try to export a vfat file larger than 4 GB, the export will fail.
- For Linux-based VMs, when exporting a VM to another system, you do not need to modify the `/etc/fstab` file.
- For Ubuntu-based VMs running some older Ubuntu releases, you may need to edit the `/boot/grub/grub.cfg` file and change the `gfxmode` parameter to `text` (for example, `set gfxmode=text`) before exporting a VM; otherwise, the new VM's console may hang on another system. You can restore the original setting in the source VM after the migration.

Prerequisites:

- You must shut down a VM before exporting it.
- Prepare the export destination:
  - If you are using a USB device, insert it into the current primary node of the system (displayed as **node*n* (primary)** on the **Physical Machines** page). Confirm that system displays the USB device. Navigate to the **Physical Machines** page. Click the node into which you inserted the device, and in the lower pane, select the **USB Device** tab. The USB device you inserted should appear in the tab's display.
  - If you are using a network-mounted folder for a Windows/CIFS share or an NFS export, create the folder in your environment where you can store the exported VM. Set full read/write permissions on the network-mounted folder to permit file transfers, or, for a Windows/CIFS share only, assign read/write permissions to a specific user on the system/domain that hosts the share. Record the URL or pathname of the NFS export or CIFS share as well as the username/password of the CIFS share, which you use when you export the VM.

  Ensure that you have enough storage for the VMs that you want to export.

In addition, Windows-based VMs require Windows-specific preparation.

To prepare for exporting a VM (Windows-based VMs only)

1. Log on to the ztC Edge system with the ztC Console.
2. On the **Virtual Machines** page, select the VM to export.
3. Click **Console** to open the console of the VM and log on to the Windows guest operating system.
4. Ensure that all volumes are labeled accurately, as summarized in Managing Windows Drive Labels.
5. Run the Windows System Preparation Tool (`Sysprep`) to prepare the guest operating system for redeployment.

To export a VM

1. Log on to the ztC Edge system with the ztC Console.
2. On the **Virtual Machines** page, select the VM that you want to export, and click **Shutdown**.

Wait for the VM to shut down. See The Virtual Machines Page.

3. With the VM selected, click **Export** to open the export wizard.
4. Select one of the following:

Note: If you have already mounted a location using the **Mount** button (as described in Mounting a USB Device or Network-mounted Folder on the ztC Edge System), the export wizard displays the mounted device URL in green. To change it, click the **Change** button.

- **Mount device via Windows Share (CIFS/SMB)**

  The export destination is a folder on a CIFS share. Enter a **Username**, **Password**, and **Repository** value. For **Repository**, enter a value in the format ***machine_URL**ShareName*** (for example, **\\192.168.1.34\MyExportVMs**).

- **Mount device via NFS**

  The export destination is a folder on a remote system, accessed through NFS. Enter a **Repository** value, which is the URL of the remote system, in the format ***nnn.nnn.nnn.nnn*** (do not include **http://** or **https://**).

- **Mount USB**

  For **USB partition list**, select a partition from the pull-down menu.

5. For **Export Path: `/mnt/ft-export:`**, type the path of the location where you want the VM to be exported and its OVF and VHD files to be stored. For example, if you want to export the VM to a new folder named `ocean1`, type `ocean1`.
6. Click **Mount**.

   If the mount succeeds, the repository appears under **Device URL** and the **Export VM** button becomes active; otherwise, an alert appears.
7. Select the volumes to include under **Boot Volume to Export** and **Data Volumes to Export**. (The boot volume is required.)
8. Click **Export VM** to export the VM.

You can monitor the **Export Status** in the **Summary** tab for the VM that you are exporting. Progress is reported as the percentage (%) completed for the whole export and for each volume. When the process is complete, the status changes to **Export completed successfully**.

To cancel the export, click **Cancel** next to the **Export progress** percentage. A dialog box opens, asking you to confirm the cancellation. Click **Yes** to cancel.

The ztC Edge system exports the VHD files (volumes) first, then it exports the OVF file. You can confirm that the process is finished when the OVF file appears in the folder.

After the export process, if you want to import or restore the OVF and VHD files on an ztC Edge system, see Importing an OVF File.

To unmount the device, see Mounting a USB Device or Network-mounted Folder on the ztC Edge System.

Troubleshooting

If necessary, use the following information to resolve problems with the export process.

To clean up after a canceled or failed export from the ztC Edge system

Remove the VM files from the export folder or create a new folder for a subsequent export.

Related Topics

Attaching a USB Device to a Virtual Machine

Creating and Migrating Virtual Machines

Managing Virtual Machine Resources

Managing the Operation of a Virtual Machine

Mounting a USB Device or Network-mounted Folder on the ztC Edge System

You can mount (or unmount) a USB device or a network-mounted folder (that is, a directory) on the ztC Edge system using the **Mount** (or **Unmount**) button on the **Virtual Machines** page. Mounting a location makes it available to the primary node at the mount point `/mnt/ft-export/`. You can then export a VM on the primary node to the mounted location, or import a VM from the mounted location to the ztC Edge system. When the export or import is finished, use the **Unmount** button to unmount the location.

(If you need to mount a USB device to access the device in the guest operating system of a VM, see Attaching a USB Device to a Virtual Machine.)

Notes:

1. You cannot unmount a mounted location that is in use. For example, you cannot unmount a location while a VM is being exported or imported.
2. The Stratus Redundant Linux software on ztC Edge systems does not support the exFAT File system. Before you mount a USB medium, format the device with NTFS. (By default, most USB media are formatted with the FAT file system, which has a limited file size of 4 GB that may be too small for most VMs.)

Prerequisite: Prepare the mount location:

- If you are using a USB device to export or import a VM, attach the device to the current primary node for the system (displayed as **node*n* (primary)** on the **Physical Machines** page). Confirm that the system displays the USB device: navigate to the **Physical Machines** page, click the node to which you attached the device, and in the lower pane, select the **USB Device** tab. The USB device you attached should appear in the tab's display.
- If you are using a network-mounted folder for a Windows/CIFS share or an NFS export, create the folder in your environment where you can store the exported VM. Set full read/write permissions on the network-mounted folder to permit file transfers, or, for a Windows/CIFS share only, assign read/write permissions to a specific user on the system/domain that hosts the share. Record the URL or pathname of the NFS export or CIFS share as well as the username/password of the CIFS share, which you use when mounting an NFS export of CIFS share.

To mount a USB device or network-mounted folder

1. On the **Virtual Machines** page, select a VM.
2. In the lower pane, click the **Mount** button.
3. Select one of the following for the mount point **/mnt/ft-export/**:
   - **Mount device via Windows Share (CIFS/SMB)**

     The mount location is a folder on a CIFS share. Enter a **Username**, **Password**, and **Repository** value. For **Repository**, enter a value in the format ***machine_URL**ShareName*** (for example, **\\192.168.1.34\MyMountLocation**).

- **Mount device via NFS**

  The mount location is a folder on a remote system accessed through NFS. For **Repository**, enter the URL of the remote system in the format ***nnn.nnn.nnn.nnn*** (do not include **http://** or **https://**).

- **Mount USB**

  For **USB partition list**, select a partition from the pull-down menu.

4. Click **Mount**.

The location is mounted on the primary node, and the **Mount** button changes to **Unmount**.

To unmount a USB device or network-mounted folder

1. On the **Virtual Machines** page, select a VM.
2. In the lower pane, click the **Unmount** button.
3. A **Confirm** dialog box appear, asking if you are sure you want to unmount the location. Click **Yes** to unmount it.

The location is unmounted, and the **Unmount** button changes to **Mount**.

Related Topics

Exporting a Virtual Machine

Managing Virtual Machines

Managing Windows Drive Labels

Label volumes in a Windows-based virtual machine to ensure that they are correctly mapped before you export the virtual .

Caution: Ensure that each volume has a unique identifiable label before running **Sysprep** (to prepare for an export). This process requires administrator privileges.

To set a label from the command prompt, type:

```
C:\>label C:c-drive
```

To list and verify all volume labels, use the **diskpart** utility:

```
C:\> diskpart

DISKPART> list volume
```

```
...
DISKPART> exit
```

After importing the virtual machine, use **Disk Manager** to reassign the drive letters. The labels you assigned before the export will help to identify the drives. For instructions on reassigning drive letters on a Windows system, search for the Microsoft Support web site.

Related Topics

Creating and Migrating Virtual Machines

Configuring Windows-based Virtual Machines

Configuring Windows-based Virtual Machines

After installing a Windows-based virtual machine, configure the additional resources and software necessary for production use, as described in:

- Updating the VirtIO Drivers (Windows-based VMs)
- Creating and Initializing a Disk (Windows-based VMs)
- Installing Applications (Windows-based VMs)

In addition, ensure that you configure the following settings:

- Change the time zone in the guest operating system to correspond to the time zone configured on the **Date and Time** preference page in the ztC Console (see Configuring Date and Time); otherwise, the VM's time zone changes whenever VMs restart or migrate. Network Time Protocol (NTP) is recommended for both the VM and the ztC Edge system.
- Disable hibernation (enabled by default in some cases) to prevent the guest operating system from going into a power-saving state.
- Configure the power button action in the guest operating system to shut down the guest (and not to hibernate it) to allow the **Shutdown** VM button in the ztC Console to work properly (see Shutting Down a Virtual Machine).
- Configure the guest operating system to generate a crash dump file if the operating system crashes. Follow the instructions in the Microsoft article, *How to generate a complete crash dump file or a kernel crash dump file by using an NMI on a Windows-based system* (Article ID: 927069). Follow the instructions in the **More Information** section.

Related Topics

Managing Virtual Machines

Updating the VirtIO Drivers (Windows-based VMs)

Update the Red Hat VirtIO drivers in your Windows-based virtual machines (VMs) to the latest supported versions, to ensure the proper operation of the VMs. For example, you should update the VirtIO drivers after upgrading the system software (Upgrading Stratus Redundant Linux Software) or after using the P2V client to migrate a VM or a physical machine (PM) to the ztC Edge system (Migrating a Physical Machine or Virtual Machine to a System).

To update the VirtIO drivers, first download them and create a VCD. Then, update the drivers in the VMs.

Notes:

- For proper operation, ensure that you download the VirtIO drivers only from the **ztC Edge Support** page, as described in the following procedure. The VirtIO ISO file on the support page contains versions of the VirtIO drivers that have been tested with the Stratus Redundant Linux software, and they are known to work. VirtIO drivers from other sources could have compatibility issues.
- When updating the VirtIO drivers, use only the **Browse my computer for the driver software** option and select the specific folder or .inf file that applies to the guest operating system. If you use the **Search automatically for updated driver software** option or select only the top level of the VirtIO VCD, Windows might automatically install an incorrect driver.
- In some cases, the guest operating system requests a restart after drivers are updated. If so, restart the guest operating system.

To download the VirtIO drivers and create a VCD

1. Download the VirtIO ISO file from the **Downloads** page at https://www.stratus.com/services-support/downloads/?tab=ztcedge.
    a. On the **Downloads** page, click **ztC Edge** (if it is not already displayed) and then select the appropriate version.

b. Scroll down to **Drivers and Tools** and then continue scrolling to **ztC Edge VirtIO Driver Update**.

c. Click the link to the appropriate file.

Ensure that you download the version of the VirtIO ISO file that matches the version of your ztC Edge system.

2. If you want to verify the integrity of the ISO image, use the associated fciv checksum file and the Microsoft File Checksum Integrity Verifier (FCIV) executable file, if you have already downloaded the Microsoft file to your system. Otherwise, use the MD5 checksum hash function.

Using fciv

Download the fciv verification file from the **Drivers and Tools** section of the **Downloads** page. Save the VirtIO Driver fciv verification file to the directory that contains the VirtIO ISO file you downloaded.

Open a command prompt window. From the directory containing the ISO, executable, and verification files, type a command similar to the following to verify the ISO image:

fciv -v -xml virtio-win-*n.n.nn*.xml

If the command succeeds (that is, it returns the message All files verified successfully), go to the next step. If the command fails, repeat the download.

Using MD5

Open a command prompt window as an administrator, and enter the following:

CertUtil -hashfile *path_to_file* MD5

The **CertUtil** command displays a message indicating whether or not it completed successfully. If the command succeeds, continue with the next step. If the command fails, repeat the download.

3. Open the ztC Console and create a VCD of the VirtIO ISO file (see Creating a Virtual CD).

To update the VirtIO drivers in a Windows-based virtual machine

1. Open the ztC Console and insert the VCD into the Windows-based VM (see Inserting a Virtual CD).
2. In the VM console window, open **Device Manager** in the guest operating system.

The method to open Device Manager varies depending on the release of the guest operating system. One method is to open the Control Panel and select **Device Manager**. Another method is to open a search window and type **Device Manager**.

3. Expand **Network adapters** and locate the **Red Hat VirtIO Ethernet Adapter**. There may be more than one adapter present depending on the number of network interfaces in your VM.

   If the **Red Hat VirtIO Ethernet Adapter** is not present, the VirtIO driver is not installed. Expand **Other devices** and locate the unknown **Ethernet Controller** device. Update the driver for this device.

   a. Right-click the **Red Hat VirtIO Ethernet Adapter** (or **Ethernet Controller**) and select **Update Driver Software**. Click **Browse my computer for the driver software**, specify the location of the VirtIO Ethernet driver (**netkvm**) for your guest operating system, and finish updating the driver. (For example, to update the driver in a Windows Server 2012 R2 guest, select the NetKVM\2k12R2\amd64**netkvm.inf** file on the VirtIO VCD.)

   b. Repeat the driver update for each additional **Red Hat VirtIO Ethernet Adapter** (or **Ethernet Controller**) device.

4. Expand **Storage controllers** and locate the **Red Hat VirtIO SCSI controller**. There may be more than one controller present depending on the number of volumes in your VM. If the **Red Hat VirtIO SCSI controller** is not present, the VirtIO driver is not installed. Locate the unknown **SCSI controller** device, and update the driver for this device:

   a. Right-click the **Red Hat VirtIO SCSI** controller (or **SCSI controller**) and select **Update Driver Software**. Click **Browse my computer for the driver software**, specify the location of the VirtIO SCSI driver (**viostor**) for your guest operating system, and finish updating the driver. (For example, to update the driver in a Windows Server 2012 R2 guest, specify the viostor\2k12R2\amd64**viostor.inf** file on the VirtIO VCD.)

   b. Repeat the driver update for each additional **Red Hat VirtIO SCSI** (or **SCSI controller**) device.

   > **Caution**: Although the device name is the **Red Hat VirtIO SCSI** controller, you must select the storage driver file that is labeled **viostor**, and not **vioscsi** (if present). Installing the **vioscsi** driver may crash your VM.

5. If applicable, restart the guest operating system to load the updated drivers.

Related Topics

Configuring Windows-based Virtual Machines

Creating and Migrating Virtual Machines

Managing the Operation of a Virtual Machine

Creating and Initializing a Disk (Windows-based VMs)

Create and initialize a disk to prepare it for partitioning into volumes in a Windows-based virtual machine.

To create and initialize a disk in a Windows-based virtual machine

1. Use the ztC Console to create a new volume on the ztC Edge system, as described in Creating a Volume in a Virtual Machine.
2. In the Windows guest operating system, open **Disk Management** or a similar utility.
3. Initialize the newly-added disk. (You may be prompted to do so automatically.)
4. Convert the disk to a dynamic disk.
5. Create one or more simple volumes on the disk.
6. Restart the Windows guest operating system.

See your Windows documentation for complete instructions.

Note: Because the Stratus Redundant Linux software already mirrors data at the physical level, volume redundancy is not required in the Windows guest operating system.

Related Topics

Opening a Virtual Machine Console Session

Configuring Windows-based Virtual Machines

Creating and Migrating Virtual Machines

Managing the Operation of a Virtual Machine

Installing Applications (Windows-based VMs)

Install an application in a Windows-based virtual machine by doing one of the following:

- Download the installation program to the guest operating system as an executable file or ISO file.
- Mount a network drive that contains the installation program.
- Create and insert a Virtual CD (VCD) that contains the installation program. See Managing Virtual CDs.

Related Topics

Opening a Virtual Machine Console Session

Configuring Windows-based Virtual Machines

Creating and Migrating Virtual Machines

Managing the Operation of a Virtual Machine

Configuring Linux-based Virtual Machines

After installing a Linux-based virtual machine, configure the additional resources and software necessary for production use, as described in:

- Creating and Initializing a Disk (Linux-based VMs)
- Installing Applications (Linux-based VMs)

In addition, ensure that you configure the following settings:

- Disable hibernation (enabled by default in some cases) to prevent the guest operating system from going into a power-saving state.
- Configure the power button action in the guest operating system to shut down the guest (and not to hibernate it) to allow the **Shutdown** VM button in the ztC Console to work properly. For the minimal server version of Ubuntu Linux, optionally install the `acpid` package to enable the **Shutdown** button. See Shutting Down a Virtual Machine.
- Install the `kexec-tools` package and configure the guest operating system to generate a crash dump file if the system crashes.
- For Ubuntu Linux guest operating systems, to prevent a problem where the VM console hangs in ztC Console, edit the `/boot/grub/grub.cfg` file and change the `gfxmode` parameter to `text` (for example, `set gfxmode=text`). If the VM console hangs before you can set the parameter, see the troubleshooting information in Opening a Virtual Machine Console Session to resolve the issue.

For more information about these settings, see your Linux documentation.

Related Topics

Managing Virtual Machines

Creating and Initializing a Disk (Linux-based VMs)

Create and initialize a disk to make it available for storing data in a Linux-based virtual machine.

To create and initialize a disk in a Linux-based virtual machine

1. In the ztC Console, create a new volume, as described in Creating a Volume in a Virtual Machine.
2. In the Linux-based virtual machine, use the volume management tool or edit files as needed to initialize and mount the volume. See your Linux documentation for complete instructions.

The disk device names for a Linux-based virtual machine are `/dev/vda` through `/dev/vdh`, not the standard `/dev/sda` through `/dev/sdh`. The ztC Edge virtual disk volumes appear in the guest operating system and are used as if they were physical disks.

Related Topics

Opening a Virtual Machine Console Session

Configuring Linux-based Virtual Machines

Creating and Migrating Virtual Machines

Managing the Operation of a Virtual Machine

Installing Applications (Linux-based VMs)

Install an application in a Linux-based virtual machine by doing one of the following:

- Download the installation package to the guest operating system as an executable file or ISO file.
- Mount a network drive that contains the installation package.
- Create and insert a Virtual CD (VCD) that contains the installation package. See Managing Virtual CDs.

Related Topics

Opening a Virtual Machine Console Session

Configuring Linux-based Virtual Machines

Creating and Migrating Virtual Machines

Managing the Operation of a Virtual Machine

Managing the Operation of a Virtual Machine

Manage the operation of a virtual machine as described in:

- Starting a Virtual Machine
- Shutting Down a Virtual Machine
- Powering Off a Virtual Machine
- Opening a Virtual Machine Console Session
- Renaming a Virtual Machine
- Removing a Virtual Machine

For additional information configuration and troubleshooting information, see Advanced Topics (Virtual Machines).

Starting a Virtual Machine

Start a virtual machine (VM) to boot the VM's guest operating system. You can also configure a starting mode for a VM, for when the ztC Edge system boots.

To start a virtual machine

1. On the **Virtual Machines** page, select a VM.
2. Click **Start** in the bottom panel.

To configure a starting mode for a virtual machine, for when the system boots

1. On the **Virtual Machines** page, select a VM.
2. Click the **Boot** tab in the bottom panel.
3. For **Auto Start Mode**, select one of the following:
    - **Last**–Return the VM to its state when the system was shutdown: if the VM was running, the VM is restarted when the system boots; if the VM was stopped, the VM is not started when the system boots.
    - **On**–Start the VM when the system boots.
    - **Off**– Do not start the VM when the system boots.
4. Click **Save**.

Related Topics

Shutting Down a Virtual Machine

Powering Off a Virtual Machine

Managing the Operation of a Virtual Machine

Shutting Down a Virtual Machine

Shut down a virtual machine (VM) to begin an orderly shutdown of the guest operating system.

Note: You can shut down a VM with guest operating system commands. Some guests allow (or can be configured to allow) you to shut down a VM using the ztC Console.

Shutting down a VM in the ztC Console is similar to pressing the power button on a physical machine, which typically results in an orderly shutdown of the operating system. In some cases, you may need to explicitly enable this feature in the guest operating system. For example:

- For any guest, verify that the power button action is set to shut down the guest operating system and not to hibernate it. If you click **Shutdown** in the ztC Console for a guest that is set to hibernate, the VM remains in a **stopping** state and never properly shuts down.
- On some guests, the power button does not shut down the system unless a user is logged on to the operating system. You may be able to update security settings to enable the power button even in the absence of a login session.
- On some minimal server versions of Ubuntu, the `acpid` package that enables the power button is not included in the default installation. You can manually install this package to enable the power button using the following command (or see the documentation for your guest operating system):

  ```
  sudo apt-get install acpid
  ```

  For versions of Ubuntu running the desktop, the ztC Console **Shutdown** button causes the VM's Ubuntu desktop to prompt you with selecting one of three icons: suspend, sleep, or shutdown. To allow the Ubuntu VM to shutdown without the desktop prompts, you must modify the `powerbtn` file.

  To modify the `powerbtn` file

  1. In the VM, edit the `/etc/acpi/events/powerbtn` file.
  2. Comment out these lines:

     ```
     event=button[ /]power
     action=/etc/acpi/powerbtn.sh
     ```

3. Add these lines:

```
event=button/power (PWR.||PBTN)
action==/sbin/poweroff
```

4. Issue the following command to restart `acpid`:

```
systemctl restart acpid
```

See the documentation for your guest operating system to configure the behavior of the system power button, thus enabling the **Shutdown** button to work in the ztC Console.

To shut down a VM in ztC Console

1. On the **Virtual Machines** page, select a VM.
2. Click **Shutdown** in the bottom panel.

   A warning message appears, asking you to confirm the shutdown. Click **Yes** to shutdown or **No** to discontinue the shutdown.

If the VM is not responding, you can also **Power Off** the VM to stop it without properly shutting down the guest operating system.

Related Topics

Starting a Virtual Machine

Powering Off a Virtual Machine

Managing the Operation of a Virtual Machine

Powering Off a Virtual Machine

Power off a virtual machine (VM) to stop it without properly shutting down guest operating system.

Caution: Use the **Power Off** command only if the **Shutdown** command or guest operating system commands fail. Powering off a VM is similar to pulling the power cord, which may result in data loss.

To power off a virtual machine

1. On the **Virtual Machines** page, select a VM.
2. Click **Power Off** in the bottom panel.

Related Topics

Starting a Virtual Machine

Shutting Down a Virtual Machine

Managing the Operation of a Virtual Machine

Advanced Topics (Virtual Machines)

Opening a Virtual Machine Console Session

Open a virtual machine (VM) console session to display the console of the guest operating system running in the VM.

The following procedure describes how to open a VM console session in the ztC Console, but you can also use a remote desktop application for this purpose.

To open a VM console session

1. On the **Virtual Machines** page, select a VM.
2. Ensure that the VM is in a running state.
3. Click **Console** () in the bottom panel.

Troubleshooting

To resolve an issue where the VM console window does not open

Ask your network administrator to open ports 6900-6999 (inclusive).

To resolve an issue where the VM console window is blank

Verify that the VM is powered on and not in the process of booting. Also, click in the console window and press any key to deactivate the screen saver.

To resolve an issue where more than one VM console window is displayed and they are behaving erratically

Close all console windows and open only one console window.

To resolve an issue where the VM console window hangs on the ztC Edge system

For Ubuntu-based VMs, the VM console hangs in the ztC Console if you do not properly set the `gfxmode` parameter. In the guest operating system, edit the `/boot/grub/grub.cfg` file and

change the `gfxmode` parameter to `text` (for example, `set gfxmode=text`).

If the console hangs before you can set the parameter, do the following:

1. Restart the VM in the ztC Console.
2. At the GRUB menu, press `e` to edit the grub command.
3. On the next screen, on the `gfxmode` line, change `$linux_gfx_mode` to `text` so the line reads:

   ```
   gfxmode text
   ```

4. Press **Ctrl-x** or **F10** to boot the guest operating system.
5. To update the setting so it persists for each boot cycle, edit the `/boot/grub/grub.cfg` file and change the `gfxmode` parameter to `text` so the line reads:

   ```
   set gfxmode=text
   ```

6. Save the `/boot/grub/grub.cfg` file.

To change the terminal type in a Linux-based VM if the console screen is unreadable

By default, the Linux operating system sets the `TERM` variable to `vt100-nav`, which is not properly supported by the `vncterm` program, the basis for the VM console in ztC Console. If you use anything other than the command line, the screen becomes unreadable. To resolve this issue, change the terminal type in the Linux guest operating system:

1. Open the `inittab` file in the guest operating system.
2. In the following line, replace `vt100-nav` with `vt100` by deleting `-nav` at the end of the line. The updated line appears as follows:

   ```
   # Run gettys in standard runlevels co:2345:respawn:/sbin/agetty xvc0
   9600 vt100
   ```

3. Save the **inittab** file.

Related Topics

Starting a Virtual Machine

Shutting Down a Virtual Machine

Managing the Operation of a Virtual Machine

Renaming a Virtual Machine

Rename a virtual machine (VM) to change its name as it appears on the **Virtual Machines** page.

If you need to change the host name of the guest operating system running in a VM, use guest operating system tools.

Prerequisite: To rename a VM, you must shut it down.

To rename a virtual machine

1. On the **Virtual Machines** page, select a VM.
2. Click **Shutdown** and wait for the VM to shut down.
3. Double-click the name of the VM.
4. Type the new name. The VM name must meet the following requirements:
   - A VM name must start with a word or a number, and the name cannot include the special characters (for example, #, %, or $).
   - A VM name cannot use hyphenated prefixes such as Zombie- or migrating-.
   - A VM name has a maximum of 85 characters.
5. Press **Enter**.

Related Topics

Removing a Virtual Machine

Creating and Migrating Virtual Machines

Managing the Operation of a Virtual Machine

Removing a Virtual Machine

Remove a virtual machine (VM) to permanently delete it and optionally delete associated volumes from the ztC Edge system.

Prerequisite: Both PMs of the ztC Edge system must be online to properly remove a VM. On the **Physical Machines** page of the ztC Console, verify that both PMs are in the **running** state and that neither PM is in maintenance mode or in the process of synchronizing.

To remove a virtual machine

1. On the **Virtual Machines** page, select a VM.
2. Click **Shutdown** in the bottom panel.
3. When the VM has stopped, click **Remove**.
4. In the **Remove Virtual Machine** dialog box, activate the check box next to volumes that you want to delete. Clear the check box for volumes to save as archives or save for attachment to another VM.

   Caution: Make sure that you select the correct VM and volumes for removal. When you click **Delete VM**, these items are permanently removed.

5. Click **Delete VM** to permanently delete the VM and any selected volumes.

Related Topics

Renaming a Virtual Machine

Creating and Migrating Virtual Machines

Managing the Operation of a Virtual Machine

Managing Virtual Machine Resources

Manage virtual machine resources to reconfigure the vCPUs, memory, storage, or network resources of an existing virtual machine.

To reconfigure virtual machine resources, use the **Reprovision Virtual Machine** wizard, as described in:

- Reprovisioning Virtual Machine Resources

To reconfigure virtual machine volumes, see the following task-specific topics:

- Creating a Volume in a Virtual Machine
- Attaching a Volume to a Virtual Machine
- Detaching a Volume from a Virtual Machine
- Removing a Volume from a Virtual Machine
- Expanding a Volume on the ztC Edge System

To recover virtual machine resources, freeing space for new volumes or virtual CDs, see:

- Recovering Virtual Machine Resources

Reprovisioning Virtual Machine Resources

Reprovision a virtual machine (VM) to change its allocation of virtual CPUs (vCPUs), memory, storage, or network resources.

Launch the **Reprovision Virtual Machine** wizard by clicking **Config** in the bottom pane of the **Virtual Machines** page. The wizard steps you through the process of reallocating resources to the VM.

Prerequisites:

- Review the prerequisites and considerations for allocating vCPUs, memory, storage, and network resources to the VM, as listed in Planning Virtual Machine Resources. For more information about storage resources, see Planning Virtual Machine Storage.
- To reprovision a VM, you must shut down the VM.

To reprovision a virtual machine

1. Open the **Virtual Machines** page (see The Virtual Machines Page).
2. Select a VM and click **Shutdown**.
3. When the VM has stopped, click **Config** to display the **Reprovision Virtual Machine** wizard.
4. On the **Name and Description** page:
    a. Type the **Name** and an optional **Description** for the VM as they will appear in the ztC Console

    The VM name must meet the following requirements:

    - A VM name must start with a word or a number, and the name cannot include the special characters (for example, #, %, or $).
    - A VM name cannot use hyphenated prefixes such as Zombie- or migrating-.
    - A VM name has a maximum of 85 characters.

    b. Click **Next**.
5. On the **vCPUs and Memory** page:
    a. Specify the number of **vCPUs** and the amount of **Memory** to assign to the VM. For more information, see Planning Virtual Machine vCPUs and Planning Virtual Machine Memory.
    b. Click **Next**.

6. On the **Volumes** page, you can:

   > **Note**:
   > You cannot modify the VM boot volume, only data volumes. However, you can detach the boot volume.

   - Click **Boot Volume** to detach the boot volume.

     > **Caution**: If you detach the boot volume, the VM becomes unbootable.

     A warning appears saying that detaching the boot value causes the VM to become unbootable. If you want to undo detaching the boot volume, click **Undo Detach**.

   - Click **Detach** to disconnect a volume from a VM and keep it for future use.
   - Click **Delete** to permanently remove a volume from the ztC Edge system.
   - Select an unattached volume from a pulldown menu (if displayed) and click **Attach**.

     You can also, if applicable, click **Add New Volume** to create a new data volume. (If the button is not visible, scroll down to the bottom of the wizard page.)

     For an unattached volume or a new volume, specify the volume's parameters:

     a. Type the **Name** of the volume.

     b. Type the **Volume Size** of the volume in gigabytes (GB). For more information about allocating storage, see and Planning Virtual Machine Storage.

     c. If applicable, click **Attach** to connect a volume to a VM.

   To continue, click **Next**.

7. On the **Networks** page, activate the check box for each shared network that you want to attach to the VM.

   For each shared network that you attach, you can also optionally:

   - Set a custom MAC address (for details, see Assigning a Specific MAC Address to a Virtual Machine).
   - Set the **State** to **Enabled** or **Disabled**, which allows you to allow or block network traffic to the selected network.

   For more information, see Planning Virtual Machine Networks. To continue, click **Next**.

8. On the **Configuration Summary** page:

   > **Caution**: Make sure that any volumes marked for removal are correct. When you click **Finish**, permanent data loss occurs on disks marked for removal.

   a. Review the configuration summary. If you need to make changes, click **Back**.

   b. To accept the VM as provisioned, click **Finish**.

9. Click **Start** to restart the VM.

10. For Windows-based VMs, if you changed the number of assigned virtual CPUs in a Windows-based VM from 1 to *n* or *n* to 1, after restarting the VM at the end of the re-provisioning process, you must shut down and restart the VM a second time. This allows the VM to correctly reconfigure itself for Symmetric Multiprocessing (SMP). The VM displays odd behavior and is not usable until it is restarted.

Related Topics

Managing Virtual Machine Resources

Planning Virtual Machine Resources

Managing Virtual Machines

Creating a Volume in a Virtual Machine

Create a volume to attach a new, blank volume to a virtual machine (VM). (You can also attach an existing, unattached volume as described in Attaching a Volume to a Virtual Machine.)

> **Prerequisite**: Before creating a volume for a VM, you must shut down the VM.

To create a new volume in a VM

1. Open the **Virtual Machines** page (see The Virtual Machines Page).
2. Select a VM and click **Shutdown**.
3. When the VM has stopped, click **Config** to display the **Reprovision Virtual Machine** wizard.
4. Click **Next** on each wizard page until the **Volumes** page is displayed. (If applicable, see Reprovisioning Virtual Machine Resources to configure additional VM resources.)
5. On the **Volumes** page, click **Add a new volume**. (If the button is not visible, scroll down to the bottom of the wizard page.)

6. Under **To Be Created**, do the following:
    a. Type the **Name** of the volume as it will appear in the ztC Console.
    b. Type the **Volume Size** of the volume to create in gigabytes (GB). For more information about allocating storage, see and Planning Virtual Machine Storage.
7. Click **Next** on each wizard page until the **Configuration Summary** page is displayed. Verify the configuration changes.
8. Click **Finish** to create the volume.
9. Start the VM and prepare the volume for use in the guest operating system, as described in:
    - Creating and Initializing a Disk (Windows-based VMs)
    - Creating and Initializing a Disk (Linux-based VMs)

Related Topics

Detaching a Volume from a Virtual Machine

Removing a Volume from a Virtual Machine

Managing Virtual Machine Resources

Planning Virtual Machine Resources

Managing Virtual Machines

Attaching a Volume to a Virtual Machine

Attach a volume to connect a currently unused volume to a virtual machine.

Note: If you attach a boot volume to a VM that already has a boot volume, the newly added volume is attached as a data volume. You might want to attach a volume in this manner to diagnose a boot problem or data corruption in another VM's boot volume. After using guest operating system tools to resolve the issue, detach the volume and reattach it to its original VM.

Prerequisite: Before attaching a volume to a virtual machine, you must shut down the virtual machine.

To attach a volume to a virtual machine

1. Ensure that the volume you want to attach is not in use by another virtual machine; otherwise, you cannot attach it. Open the **Volumes** page, locate the volume, and ensure that the value in the **Used By** column is **None**.
2. Open the **Virtual Machines** page (see The Virtual Machines Page).
3. Select a VM and click **Shutdown**.
4. When the VM has stopped, click **Config** to display the **Reprovision Virtual Machine** wizard.
5. Click **Next** on each wizard page until the **Volumes** page is displayed. (If applicable, see Reprovisioning Virtual Machine Resources to configure additional VM resources.)
6. On the **Volumes** page, locate the pulldown menu next to the **Add a new volume** button. Select an unattached volume from the pulldown menu and click **Attach**.

   (If the pulldown menu is not visible, scroll down to the bottom of the wizard page. The pulldown menu is displayed only if there are unattached volumes on the ztC Edge system.)
7. Click **Next** on each wizard page until the **Configuration Summary** page is displayed. Verify the configuration changes.
8. Click **Finish** to attach the selected volume.

Related Topics

Creating a Volume in a Virtual Machine

Detaching a Volume from a Virtual Machine

Removing a Volume from a Virtual Machine

Managing Virtual Machine Resources

Planning Virtual Machine Resources

Managing Virtual Machines

Detaching a Volume from a Virtual Machine

Detach a volume to disconnect it from a virtual machine and keep it for future use, or attach it to another virtual machine as described in Attaching a Volume to a Virtual Machine. (You can also permanently delete the volume from the ztC Edge system, as described in Removing a Volume from a Virtual Machine.)

Note: If you detach a boot volume from a VM, you cannot boot the VM; however, you might want to detach the boot volume to diagnose a boot problem or data corruption in the volume. You can temporarily attach the boot volume to another VM as a data volume, as described in Attaching a Volume to a Virtual Machine. After using guest operating system tools to resolve the issue, detach the volume and reattach it to its original VM.

Prerequisite: Before detaching a volume from a virtual machine, you must shut down the virtual machine.

To detach a volume from a virtual machine

1. Open the **Virtual Machines** page (see The Virtual Machines Page).
2. Select a VM and click **Shutdown**.
3. When the VM has stopped, click **Config** to display the **Reprovision Virtual Machine** wizard.
4. Click **Next** on each wizard page until the **Volumes** page is displayed. (If applicable, see Reprovisioning Virtual Machine Resources to configure additional VM resources.)
5. On the **Volumes** page, locate the volume to detach. (If the volume is not visible, scroll down on the wizard page.)
6. Click **Detach** beside the volume name to mark the volume for detachment.

   Caution: Be careful to mark the correct volume to detach, avoiding any volumes that are currently in use.
7. Click **Next** on each wizard page until the **Configuration Summary** page is displayed. Verify the configuration changes.
8. Click **Finish** to detach the selected volume.

Related Topics

Attaching a Volume to a Virtual Machine

Removing a Volume from a Virtual Machine

Managing Virtual Machine Resources

Planning Virtual Machine Resources

Managing Virtual Machines

Removing a Volume from a Virtual Machine

Remove a virtual machine (VM) volume to permanently delete it from the ztC Edge system. (You can also detach a volume from the VM but keep it for future use, as described in Detaching a Volume from a Virtual Machine.)

Prerequisite: Before removing a volume attached to a virtual machine, you must shut down the virtual machine.

To remove a volume that is attached to a virtual machine

1. Open the **Virtual Machines** page (see The Virtual Machines Page).
2. Select a VM and click **Shutdown**.
3. When the VM has stopped, click **Config** to display the **Reprovision Virtual Machine** wizard.
4. Click **Next** on each wizard page until the **Volumes** page is displayed. (If applicable, see Reprovisioning Virtual Machine Resources to configure additional VM resources.)
5. On the **Volumes** page, locate the volume to delete. (If the volume is not visible, scroll down on the wizard page.)
6. Click **Delete** beside the volume name to mark the volume for deletion.

   Caution: Be careful to mark the correct volume to remove, avoiding any volumes that are currently in use.
7. Click **Next** on each wizard page until the **Configuration Summary** page is displayed. Verify the configuration changes.
8. Click **Finish** to permanently delete the selected volume.

To remove an unattached volume

> **Caution**: Before removing a volume, ensure that it is no longer needed by other administrators.

1. Open the **Volumes** page.
2. Select an unattached volume. (The **Used By** column must read ***None***, otherwise, the **Remove** button is not displayed.)
3. Click **Remove**.

Related Topics

Detaching a Volume from a Virtual Machine

Attaching a Volume to a Virtual Machine

Managing Virtual Machine Resources

Planning Virtual Machine Resources

Managing Virtual Machines

Renaming a Volume on the ztC Edge System

Rename a volume on the ztC Edge system to change its name as it appears on the **Volumes** page.

If you need to change the name of a disk or volume in the guest operating system running in a virtual machine, use guest operating system tools.

To rename a volume on the ztC Edge system

1. Locate the volume on the **Volumes** page.
2. Double-click the name of the volume.
3. Specify the new name and press **Enter**.

Related Topics

Creating a Volume in a Virtual Machine

Detaching a Volume from a Virtual Machine

Removing a Volume from a Virtual Machine

Managing Virtual Machine Resources

Planning Virtual Machine Resources

Managing Virtual Machines

Expanding a Volume on the ztC Edge System

Expand a virtual machine (VM) volume to allocate more space for programs and data in the guest operating system.

You can expand a volume, but you cannot reduce the size of a volume. Use the following procedure to expand a volume only when the VM is stopped.

Prerequisites:

- You must shut down the VM before expanding a volume that it contains.
- Ensure that both PMs of the ztC Edge system are online; otherwise, the system cannot properly expand a volume.

To expand a volume

1. On the **Physical Machines** page (see The Physical Machines Page) of a system licensed for two nodes, verify that both PMs are in the **running** state and that neither PM is in maintenance mode or in the process of synchronizing.
2. On the **Virtual Machines** page (see The Virtual Machines Page), select the VM that contains the volume that you want to expand. Ensure that the VM is **stopped**.
3. In the bottom pane, click the **Volumes** tab and select the volume that you want to expand. In the **Action** column, click **Expand Volume**.
4. Next to **Expand By**, type the amount of storage space to add to the volume (in gigabytes (GB)). When you type the number, the dialog box displays the **Expanded Volume Size** that will result if you complete the operation.

> **Note**: Consider the **Expand By** entry carefully, because after expanding a volume, you cannot undo the change or reduce the size of the volume; you can only expand the volume further.

5. Click **Expand Volume** to commit the change and expand the volume. The dialog box displays the expansion progress and automatically closes when the operation is complete.

Related Topics

Creating a Volume in a Virtual Machine

Detaching a Volume from a Virtual Machine

Removing a Volume from a Virtual Machine

Managing Virtual Machine Resources

Planning Virtual Machine Resources

Managing Virtual Machines

Recovering Virtual Machine Resources

To conserve storage space, remove VM resources when they are no longer needed. You may also need to immediately recover storage space when there is insufficient space for certain tasks, such as creating a volume or VCD.

To recover storage space, remove unused resources as described in the following topics:

- Removing a Virtual Machine
- Removing a Volume from a Virtual Machine
- Removing a Virtual CD

Related Topics

Managing Virtual Machine Resources

Planning Virtual Machine Resources

Managing Virtual Machines

Managing Virtual CDs

Create and manage virtual CDs (VCDs) to make software installation media available to the virtual machines on your ztC Edge system in ISO format.

A VCD is a read-only ISO image file that resides on a storage device of the ztC Edge system. Use the **Virtual CD Creation Wizard** (in ztC Console) to upload an existing ISO file, as described in Creating a Virtual CD.

After you create a VCD, you can boot from it to install a Windows or Linux guest operating system, or start a VM from a bootable recovery VCD. You can also insert a VCD into a running VM to install software applications.

You manage VCDs as described in:

- Creating a Virtual CD
- Inserting a Virtual CD
- Ejecting a Virtual CD
- Booting from a Virtual CD
- Renaming a Virtual CD
- Removing a Virtual CD

Users who are assigned the role **Administrator** or **Platform Manager** can perform all VCD tasks. Users who are assigned the role **VM Manager** can perform all VCD tasks, except rename a VCD. (For information on assigning these roles, see Managing Local User Accounts.)

Creating a Virtual CD

Create a virtual CD (VCD) to make software installation media available to the virtual machines (VM) on your ztC Edge system.

To create a VCD, use the **Virtual CD Creation Wizard** to upload or copy an ISO file to a storage device on the ztC Edge system. Thereafter, you can boot from it (see Booting from a Virtual CD) to install a guest operating system or start a VM from a bootable recovery VCD. You can also insert a VCD into a running VM (see Inserting a Virtual CD) to install software applications.

Notes:

1. Unless you use a VCD on a regular basis, remove it when it is no longer needed.
2. If you create a bootable VCD for installation, it must be a single CD or DVD. Multiple CDs or DVDs are not supported.

To create a VCD

1. If necessary, create ISO files of any physical media for which you will create VCDs.
2. Open the **Virtual CDs** page in the ztC Console.
3. Click **Create VCD** to open the **Virtual CD Creation Wizard**.
4. Type a name for the VCD.
5. Select a source for the VCD:
   - **Upload ISO file** uploads a file from your system running the ztC Console. Click **Browse**, select the ISO file on your system, and click **Open**.
   - **Copy CD ISO from network source** copies the file from a Web URL. Specify the URL of the ISO file.
6. Click **Finish** to upload or copy the ISO file from the specified source.

   The **Virtual CD Creation Wizard** displays progress of the upload.

You can determine the status of a VCD by checking the **State** column on the **Virtual CDs** page:

- A syncing icon ( ) indicates that the VCD is still being created.
- A broken icon ( ) indicates that the VCD creation failed. Remove the VCD and try creating it again.
- A normal icon ( ) indicates that the transfer is complete and that the VCD is ready to use.

Related Topics

Inserting a Virtual CD

Ejecting a Virtual CD

Managing Virtual CDs

Creating and Migrating Virtual Machines

Inserting a Virtual CD

Insert a virtual CD (VCD) in a virtual machine (VM) to access installation media when installing applications in a guest operating system. (To attach a USB device, see Attaching a USB Device to a Virtual Machine. To boot a virtual machine from a VCD, see Booting from a Virtual CD.)

Caution: When you insert a VCD into a running VM, it prevents the Stratus Redundant Linux software from migrating the VM to a different physical machine in the event of a failure. To restore redundancy, unmount and eject the VCD as soon as you finish using it.

Note: By default, VCDs are enabled for insertion in VMs. To change this configuration, see Configuring VM Devices.

To connect a VCD to a VM

1. If necessary, create a VCD (see Creating a Virtual CD) for the software installation media you need to access.
2. On the **Virtual Machines** page, select a VM.
3. In the bottom pane, click the **CD Drives & USB Devices** tab.
4. To select a VCD, click **Insert a CD** and select a VCD. Use the pulldown menu, if it exists.

When the system has inserted the VCD, its name appears to the right of **CD-ROM**.

Related Topics

Creating a Virtual CD

Ejecting a Virtual CD

Booting from a Virtual CD

Managing Virtual CDs

Ejecting a Virtual CD

Eject a virtual CD (VCD) to disconnect it from a virtual machine (VM). Ejecting a VCD allows you to insert another VCD into the VM. It also makes the VCD available for inserting into another VM.

To eject a VCD from a VM

1. Unmount the VCD from the guest operating system to ensure that it is not in use.
2. On the **Virtual Machines** page, select a VM.
3. Click the **CD Drives & USB Devices** tab in the lower frame.
4. On the **CD Drives** tab, click **Eject CD**.

Related Topics

Creating a Virtual CD

Inserting a Virtual CD

Booting from a Virtual CD

Managing Virtual CDs

Booting from a Virtual CD

Boot a virtual machine from a virtual CD (VCD) to install a guest operating system or to perform maintenance.

Before booting from a VCD, you must shut down the virtual machine.

To boot a virtual machine from a VCD

1. If necessary, create a VCD from a bootable CD/DVD (see Creating a Virtual CD).
2. On the **Virtual Machines** page, select a virtual machine.
3. If the virtual machine is running, click **Shutdown**.
4. When the virtual machine status shows **stopped**, click **Boot from CD** in the lower pane.
5. Select the bootable VCD, then click **Boot**.

Note: A Windows-based virtual machine booted from a VCD boots as a hardware virtual machine (HVM), and it can access only the first three disk volumes.

Related Topics

Creating a Virtual CD

Inserting a Virtual CD

Ejecting a Virtual CD

Managing Virtual CDs

Creating and Migrating Virtual Machines

Managing the Operation of a Virtual Machine

Renaming a Virtual CD

Rename a virtual CD (VCD) to change its name as it appears on the **Virtual CDs** page.

To rename a VCD

1. Locate the VCD on the **Virtual CDs** page.
2. Double-click the name of the VCD.
3. Specify the new name and press **Enter**.

Related Topics

Removing a Virtual CD

Inserting a Virtual CD

Ejecting a Virtual CD

Creating a Virtual CD

Managing Virtual CDs

Removing a Virtual CD

Remove a virtual CD (VCD) to permanently delete it from the ztC Edge system.

To remove a VCD

1. In the ztC Console, click **Virtual CDs**.
2. Locate the VCD you want to remove in the list.
3. Ensure that the **Can Remove** column displays **Yes** for the VCD. If the value is **No**, the VCD is

currently in use.

4. Select the VCD and click **Remove** in the lower panel.

Related Topics

Renaming a Virtual CD

Inserting a Virtual CD

Ejecting a Virtual CD

Creating a Virtual CD

Managing Virtual CDs

Advanced Topics (Virtual Machines)

The following topics describe procedures and information for advanced users:

- Assigning a Specific MAC Address to a Virtual Machine
- Selecting a Preferred PM for a Virtual Machine
- Forcing a VM to Boot
- Configuring the Boot Sequence for Virtual Machines
- Resetting MTBF for a Failed Virtual Machine
- Attaching a USB Device to a Virtual Machine

To manage the operation of a virtual machine, see Managing the Operation of a Virtual Machine.

Assigning a Specific MAC Address to a Virtual Machine

Assign a specific Media Access Control (MAC) address to a virtual machine (VM) if you need to override its default MAC address.

Warnings:

1. By default, the Stratus Redundant Linux software automatically assigns MAC addresses to the VMs. Do not override the default settings unless you have specific requirements (for example, to support software applications that are licensed on a MAC-address basis).
2. If you change the **Static System IP** address, any MAC addresses automatically assigned to the VMs will change when the VMs reboot, because the Stratus Redundant Linux software generates MAC addresses for the VMs based on the system IP address. To prevent changes to the MAC address for a VM, set a persistent MAC address as described in the following procedure. Contact your network administrator to generate a valid MAC address for your environment, and remember to update any firewall rules based on the new MAC address.

Prerequisite: Before modifying the MAC address for a virtual machine, you must shut down the VM.

To assign a specific MAC address to a VM

1. Open the **Virtual Machines** page (see The Virtual Machines Page).
2. Select a VM and click **Shutdown**.
3. When the VM has stopped, click **Config** to display the **Reprovision Virtual Machine** wizard.
4. Click **Next** on each wizard page until the **Networks** page is displayed. (If applicable, see Reprovisioning Virtual Machine Resources to configure additional VM resources.)
5. On the **Networks** page, locate the network to modify and make a note of the current MAC address in case you need to restore it.
6. Type the new address in the **MAC address** column, or leave the text area blank to allow the Stratus Redundant Linux software to automatically assign the MAC address.
7. Click **Finish**.

Related Topics

Advanced Topics (Virtual Machines)

Managing Virtual Machine Resources

Managing the Operation of a Virtual Machine

Selecting a Preferred PM for a Virtual Machine

On systems licensed for two nodes, select a preferred physical machine to ensure that a virtual machine runs on a particular physical machine in the ztC Edge system.

> **Note**: By default, the system automatically balances the load of virtual machines over the two physical machines. Do not modify this setting unless you have specific load balancing requirements .

To select a preferred physical machine

1. On the **Virtual Machines** page, select a virtual machine.
2. In the bottom pane, click the **Load Balance** tab.
3. Choose your preference from the pulldown list and click **Save**.

Related Topics

Advanced Topics (Virtual Machines)

Managing the Operation of a Virtual Machine

Forcing a VM to Boot

You can force a VM to boot using the **Force Boot** button on the VIRTUAL MACHINES page. However, the **Force Boot** button is active only when the ztC Console reports that the partner node is powered off or otherwise unreachable. When you use **Force Boot** to bring a VM online, you manually by-pass the system's safety checks to protect data, so you must use **Force Boot** with extreme caution and with full understanding of the conditions and consequences of using it.

> **Caution**: Before using **Force Boot**, read this entire topic and consult with your authorized Stratus service representative. The service representative can review your system, including the date of the last volume synchronization, and can then discuss with you the full impact of using **Force Boot**. Then, you can decide, with your service representative, whether or not to force a VM to boot.

When you force a VM online with **Force Boot**, you select a node (that is, the node that is reachable) on which to force the VM to boot. All data on that node is marked as valid, regardless of the actual condition of the data (for example, the data's state, the last synchronization, the condition of the volume, etc.).

During the **Force Boot** process, the VM's volumes are tagged with the date and time that the force-boot process was initiated. The VM's AX components (that is, the VM's AX pair) use the data on the VM's volumes and communicate the status of that data to determine which AX contains the up-to-date volume information. The **Force Boot** process overrides the built-in logic that protects a VM from running in a split-brain condition. If the AX pair cannot communicate, a split-brain condition occurs and damages data integrity (for information on the split-brain condition, see Creating an ALSR Configuration).

Warnings: Do not use **Force Boot** in the following situations:

- One or more volumes is the target of an unfinished mirror copy on the node where you will perform **Force Boot**.
- A target of an unfinished mirror copy is not good and will not be available even with **Force Boot**.
- The volumes are not synchronized. The following two situations are examples:
  - Both of the VM's AXs must have access to all of the VM's data volumes.
  - On a system with multiple volumes, the VM needs both AXs to be running in order for the VM to have access to all of its volumes because each node has a green-checked copy of a different volume, and the volume's mirror copy on the opposite node is not green-checked.
- Both nodes are required because multiple VMs are degraded, yet are green-checked on opposite nodes (for example, Node0 has a good boot volume but a bad data volume, while Node1 has a bad boot volume but a good data volume).
- The system is licensed for one node.

If you perform a **Force Boot** on a system with outdated volumes, contact your authorized Stratus service representative immediately. If both nodes are powered on and have started to synchronize data, the system uses data from the VM that you forced to boot, and you cannot recover the data on the node that was unreachable.

In some circumstances, however, you might be able to recover data after you use **Force Boot** on a system with outdated volumes:

- If the unreachable node is still powered off, do not power it on.
- If the unreachable node was powered off before you clicked **Force Boot**, then the VM's AX on the powered-off node is preserved and you can reverse the **Force Boot** without data loss under the following conditions:

- The VM that you forced to boot does not have new data (that is, the VM has not been put in production).
- Before you forced the VM to boot, the VM's AX on the unreachable node did not exchange status with the AX of the VM that you will force to boot.
- The issue preventing the VM's AX on the unreachable node from booting is resolved.
- All VM data between the two nodes is accurately synchronized. The system has no VMs where, of each VM's two AX components, the data of the VM's AX on one node is in a different state from the data of the VM's AX on the other node.

If your system meets all of the conditions above, contact your authorized Stratus service representative to advise you on a recovery process.

If you have decided to force a VM to boot, be sure to prepare for it by performing the prerequisite procedures.

Prerequisites:

- Manually check all volumes to ensure that you can safely override them. For example, the volume state should be green-checked, and disk synchronization should be finished.
- Determine if both AX components of the VM can communicate and can allow the system processes to determine the state of each volume. To prevent a split-brain condition, you must ensure that the two AX components of the VM can communicate status and can determine which AX has good data and good boot volumes.
- Ensure that the system is licensed for two nodes.
- Contact your authorized Stratus service representative.

To force a VM to boot

After you have consulted with your authorized Stratus service representative, and you have decided to force a VM to boot, perform the following procedure. In the examples, node0 is offline, node1 is the primary, and VM-1 is stopped.

1. In the ztC Console of a system licensed for two nodes, click **Virtual Machines** in the left panel.
2. Navigate to the **Virtual Machines** page.
3. On the **Virtual Machines** page, select the VM that is stopped and that you want to force to boot (for example, VM-1).
4. In the bottom panel, click the **Start** button.

The VM begins to boot. It continues booting until the time-out limit is reached, possibly as long as 5 minutes. When the time-out limit is reached, the **Force Boot** button becomes active.

5. To force the VM to boot, click **Force Boot**.

   A warning appears, asking you if you are positive that you know which node has the most up-to-date VM data. The warning also tells you to be aware that data loss is possible. In addition, a message tells you the node on which you can force the VM to boot.

   Caution: If you select the wrong node during **Force Boot**, data is damaged.

   You must type the node (node0 or node1) as indicated in the message. The following message is an example:

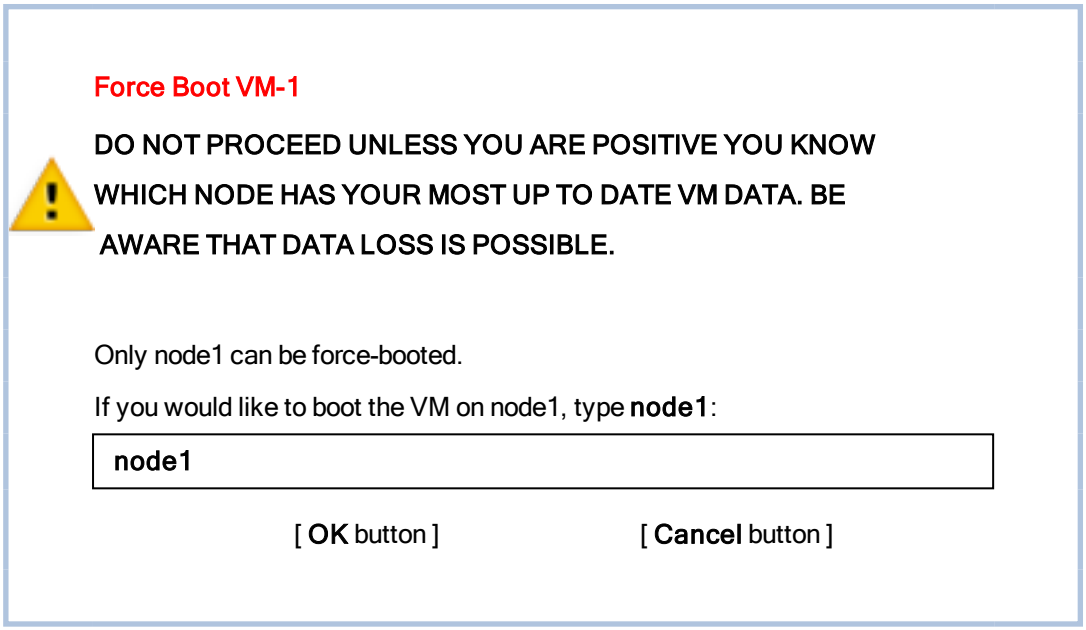

6. Click **OK** to force the node (for example, node1) to boot. (Click **Cancel** to cancel the procedure.) As the force-boot process begins and continues, additional confirmation messages appear before the VM starts and the data is marked as valid to the system.

   The VM begins to run. On the **Virtual Machines** page, the VM is listed with a warning because the node (for example, node0) is still offline.

   Once the secondary node is brought back in to the system, all data synchronizes from the node running the VM. In this example, all data synchronizes from node1 to node0.

Related Topics

Advanced Topics (Virtual Machines)

Managing the Operation of a Virtual Machine

Configuring the Boot Sequence for Virtual Machines

Configure the boot sequence of virtual machines to set the order in which guest operating systems and applications are started on the ztC Edge system.

Determine the required boot sequence, then configure the boot settings for each virtual machine accordingly.

To set the boot sequence for a virtual machine

1. On the **Virtual Machines** page, select a virtual machine.
2. In the bottom pane, click the **Boot Sequence** tab.
3. Configure the boot settings, as described below.
4. Click **Save**.

The boot settings are as follows:

- **Priority Group** enables users to specify the order in which virtual machines boot after powering on the ztC Edge system or after a failover, which requires restarting virtual machines. Some business solutions require specific virtual machines to be running before starting others. Group **1** is the highest priority and **none** is the lowest. The Stratus Redundant Linux software waits for the **OS and Application Start Time** to elapse before starting virtual machines in the next priority group.

  Boot sequence example:

VM	Priority Group	OS and Application Start Time
DNS	1	2 mins
App	2	30 secs
DB	2	10 mins
Web	3	0

1. ztC Edge boots the DNS VM.
2. 2 minutes after the DNS VM is started, ztC Edge starts the App and DB servers in group 2.
3. 10 minutes after the DB VM is started, ztC Edge starts the Web VM in group 3.

- **OS and Application Start Time** should be set to the time it takes from starting the virtual machine until the guest operating system and applications are fully functional.

Related Topics

Advanced Topics (Virtual Machines)

Managing the Operation of a Virtual Machine

Resetting MTBF for a Failed Virtual Machine

Reset the mean time between failure (MTBF) counter for a virtual machine to attempt to restart a failed virtual machine.

If a virtual machine's guest OS crashes, ztC Edge automatically restarts it, unless it has fallen below its MTBF threshold. If the virtual machine is below the MTBF threshold, ztC Edge leaves it in the crashed state. If necessary, you can reset the MTBF counter and restart the virtual machine.

Caution: Do not reset the MTBF counter unless instructed to do so by your authorized Stratus service representative, as doing so may affect the continuous uptime of your system.

Notes:

1. The **Reset Device** button is displayed only if the virtual machine falls below its MBTF threshold.
2. The **Clear MTBF** button is displayed only if the system software supporting a VM on one physical machine falls below its MBTF threshold.

To reset the MTBF counter for a virtual machine

1. On the **Virtual Machines** page, select a virtual machine.
2. Click **Reset Device**.

If the system software supporting a VM on one physical machine fails too often, perform the steps below to reset its MTBF counter.

To reset the MTBF counter for a VM on one physical machine

1. On the **Virtual Machines** page, select a virtual machine.
2. Click **Clear MTBF**.

Related Topics

Advanced Topics (Virtual Machines)

Managing the Operation of a Virtual Machine

Creating a Diagnostic File

Attaching a USB Device to a Virtual Machine

Attach a USB device to a virtual machine (VM) in order to enable the VM to use the device. A USB device may be needed, for example, when a USB-based license is required to install an application in a guest operating system. When you no longer need the USB device, detach it.

(If you need to mount a USB device on the ztC Edge system to use the device for exporting or importing VMs, see Mounting a USB Device or Network-mounted Folder on the ztC Edge System.)

Notes:

1. You can attach only supported USB devices to a guest operating system. ztC 100i systems provide two USB 3.0 ports, but do not provide a USB 3.1 port. On ztC 100i systems, USB 3.1 devices and USB 3.0 devices are supported when inserted into the USB 3.0 port.
2. Do not insert a USB 3.0 (or higher) device into a USB 3.0 host port. The system does not support a USB 3.0 (or higher) device inserted into a USB 3.0 port. You can, though, insert a USB 3.0 (or higher) device into a USB 2.0 host port.
3. The VM must be running in order for you to attach a USB device to it.
4. By default, USB devices are enabled for attachment to VMs. To change this configuration, see Configuring VM Devices.

To attach a USB device to a VM

1. Insert the USB device into the primary (active) node for the VM.

The **Virtual Machines** page displays the primary node for each VM as the **Current PM**. (This node may be different from the current primary node for the ztC Edge system, as displayed on the **Physical Machines** page.)

Confirm that the system displays the USB device. Navigate to the **Physical Machines** page. Click the node into which you inserted the device, and in the lower pane, select the **USB Device** tab. The USB device you inserted should appear in the tab's display.

2. On the **Virtual Machines** page, select a VM.
3. In the bottom pane, click the **CD Drives & USB Devices** tab.
4. On the **USB** line of the **CD Drives & USB Devices** tab, select a USB device from the pull-down menu.
5. Click **Attach a USB** to attach the USB device to the VM.
6. A **Confirm** dialog box appear, asking if you are sure you want to attach the device and displaying a warning that the guest goes simplex while the USB device is in use. Click **Yes** to attach the device.

   After the system attaches the USB device to the VM, the USB device name appears in the list of USB devices on the **CD Drives & USB Devices** tab for the VM.

To detach a USB device from a VM

1. On the **Virtual Machines** page, select the VM to which the USB device is attached.
2. In the bottom pane, click the **CD Drives & USB Devices** tab.
3. On the **USB** line of the **CD Drives & USB Devices** tab, click **Detach USB device**. If necessary, select the USB device from the pull-down menu.
4. A **Confirm** dialog box appear, asking if you are sure you want to detach the device. Click **Yes** to detach the device.

   After the system detaches the USB device to the VM, the USB device name no longer appears in the list of USB devices on the **CD Drives & USB Devices** tab for the VM.

Related Topics

Managing Virtual Machines

7

Chapter 7: Maintaining Physical Machines

You can maintain physical machines (PMs), or nodes, in a ztC Edge system by replacing or recovering them.

To replace a failed PM, use one of the following procedures:

- Replacing Physical Machines (Automated) (Recommended)

  Describes how to replace a failed PM with the automated node replacement process. This help topic supplements the information in *ztC Edge 100i Systems: Node Replacement* (R002Z), which is included with each replacement node.

- Replacing Physical Machines (Manual)

  Describes how to replace a failed PM with the user-initiated replacement process, which you start and monitor from the ztC Console. Avoid using this user-initiated procedure unless specifically instructed by your authorized Stratus service representative.

To recover the system software on a failed PM instead of replacing the PM hardware, see Recovering a Failed Physical Machine (Manual).

Replacing Physical Machines (Automated)

This topic describes how to replace a failed physical machine (PM), or node, in a dual-node ztC Edge system with the automated node replacement process. It supplements the information in *ztC Edge 100i Systems: Node Replacement* (R002Z).

You replace a node in a dual-node system while the system is running.

Prerequisite: To request a replacement ztC Edge node, log on to the **Stratus Customer Service Portal**, expand **Customer Support**, and click **Add Issue**. When creating the issue, please have the following information ready:

- **Asset ID**–Locate the **Asset ID** for your system in the masthead of the ztC Console window.
- Diagnostic file–Generate and download a diagnostic file on the **Support Logs** page of the ztC Console, as described in Creating a Diagnostic File. Attach the diagnostic file to the issue that you add in the Service Portal.

A customer service representative will contact you to diagnose the issue and provide a replacement node, if necessary.

To replace a node in a ztC Edge system

1. Locate the node to replace. The faulted node is either powered off (automatically) or powered on with the SYS LED off or solid green (not healthy). If the node is already powered off, skip to step 3.
2. If the faulted node is still powered on, open the ztC Console to resolve any issues blocking shutdown. For example, a failed network connection on the healthy node can cause a dependency on the faulted node. Resolve any issues and shut down the faulted node.
3. Disconnect the power cable from the faulted node, then disconnect the network cables and remove the node from system.
4. Add the replacement node to the system. Reconnect the network cables, then reconnect power to automatically power on the node. The node replacement is complete. The system begins to synchronize with no user input required.
5. After 20 minutes, the SYS LED cycles from off to solid green to show that the software on the replacement node is starting. After another 15 minutes, the SYS LED starts flashing to show that the system is healthy.
6. Log on to the ztC Console to confirm the system health. The virtual machines may continue to synchronize for hours. After synchronization completes successfully, the **Dashboard** should display green check marks with no outstanding issues.

Related Topics

Maintenance Mode

Maintaining Physical Machines

The ztC Console

Physical Machines and Virtual Machines

The Physical Machines Page

Replacing Physical Machines (Manual)

Caution: If you need to recover or replace a PM in a ztC Edge system, use the instructions in *ztC Edge 100i Systems: Node Replacement* (R002Z). (If needed, see Replacing Physical Machines (Automated) for additional details.) Avoid using the manual procedure described in this topic unless specifically instructed by your authorized Stratus service representative.

You replace a physical machine (PM), or node, of a dual-node ztC Edge system while the system is running. (If you need to recover the system software on a failed PM instead of replacing the PM hardware, see Recovering a Failed Physical Machine (Manual).)

When you remove and replace a PM, the system completely erases all of the disks in the replacement PM in preparation for a full installation of the Stratus Redundant Linux system software. To install the software, you can allow the system to automatically boot the replacement node from a temporary Preboot Execution Environment (PXE) server on the primary PM. As long as each PM contains a full copy of the most recently installed software kit (as displayed on the **Upgrade Kits** page of the ztC Console), either PM can initiate the replacement of its partner PM with PXE boot installation. If needed, you can also manually boot the replacement node from USB installation media.

Use one of the following procedures based on the media you want to use for the installation, either **PXE** or **USB** installation.

Caution: The replacement procedure deletes any software installed in the host operating system of the PM and all PM configuration information entered before the replacement. After you complete this procedure, you must manually re-install all of your host-level software and reconfigure the PM to match your original settings.

Caution: To prevent data loss, if the system log indicates that manual intervention is necessary to assemble a disk mirror, contact your authorized Stratus service representative for assistance. You may lose valuable data if you force a resynchronization and overwrite the most recent disk in the mirror.

Prerequisite: To request a replacement ztC Edge node, log on to the **Stratus Customer Service Portal**, expand **Customer Support**, and click **Add Issue**. When creating the issue, please have the following information ready:

- **Asset ID**–Locate the **Asset ID** for your system in the masthead of the ztC Console window.
- Diagnostic file–Generate and download a diagnostic file on the **Support Logs** page of the ztC Console, as described in Creating a Diagnostic File. Attach the diagnostic file to the issue that you add in the Service Portal.

A customer service representative will contact you to diagnose the issue and provide a replacement node, if necessary.

Prerequisites: If you want to use a USB medium to install the system software on the replacement PM:

- Create a bootable USB medium as described in Creating a USB Medium with System Software.

  When creating the USB medium, ensure that it contains the most recently installed upgrade kit. For example, if the release shown in the masthead of the ztC Console window is version `1.2.0-550`, where 550 is the build number, the kit you select to create the USB medium on the **Upgrade Kits** page must also be version `1.2.0-550`. If the system detects a different build on the replacement PM, it automatically restarts the replacement process, initializes all data on the replacement PM, and uses PXE boot installation to reinstall the most recently installed software kit on the PM with no user interaction.
- Connect a keyboard and monitor to the replacement PM to monitor the installation process and specify settings.

To remove and replace a failed PM (with PXE boot installation)

Use the following procedure to replace a failed PM and reinstall the system software by using PXE

boot installation from the software kit on the primary PM.

1. In the ztC Console, click **Physical Machines** in the left-hand navigation panel.
2. Select the appropriate PM (node0 or node1) and then click **Work On**, which changes the PM's **Overall State** to **Maintenance Mode** and the **Activity** state to **running (in Maintenance)**.
3. After the PM displays **running (in Maintenance)**, click **Recover**.
4. When prompted to select the type of repair, click **PXE PM Replace - Initialize All Disks**.

   

   Caution: Selecting **PXE PM Replace - Initialize All Disks** deletes all data on the replacement PM.

5. Select one of the following PXE Settings:

   - **Only respond to PXE requests from the current partner node.**

     Waits for a PXE boot request from the MAC address of the current partner node. Select this option if you are recovering the existing PM by completely wiping and reinstalling it. This process deletes all data on the PM, but restores its current network configuration.

   - **Only respond to PXE requests from the following MAC address.**

     Waits for a PXE boot request from the MAC address that you specify. Select this option if you are replacing the PM with a new PM. Enter the MAC address of the specific network adapter that will initiate PXE boot.

   - **Accept PXE requests from any system on priv0.**

     Waits for a PXE boot request from priv0, the private network that connects the two ztC Edge nodes. Select this option if you are replacing the PM with a new PM, but you do not know the MAC address for the new PM.

6. Click **Continue** to begin the replacement process. The system shuts down and powers off the PM.
7. After the PM is powered off, install the replacement PM, if applicable:
   a. Disconnect and remove the old PM, and then install the replacement PM.
   b. Reconnect the network cables to their original ports, and then reconnect power.
8. If the PM does not automatically power on, press the power button.

9. The replacement process continues with no user interaction, as follows:
   - The replacement PM begins to boot from a PXE server that temporarily runs on the primary node.
   - The system automatically deletes all of the data on disks in the replacement PM.
   - The replacement PM reboots again and automatically starts the system software installation, which runs from a copy of the installation kit on the primary node.

   You do not need to monitor the progress of the software installation or respond to prompts at the physical console of the replacement PM. The replacement process is automated, and it is normal for the PM to display a blank screen for a long period of time during the software installation.
10. When the software installation is complete, the replacement PM reboots from the newly installed system software.

    Note: After the system software installation, the replacement PM may take up to 20 minutes to join the system and appear in the ztC Console.
11. As the replacement PM joins the system, you can view its activity on the **Physical Machines** page of the ztC Console. The **Activity** column displays the PM as **(in Maintenance)**, and then as **running** after the replacement is complete. The PM automatically exits maintenance mode and begins load balancing the VMs on the system.
12. If applicable, manually reinstall applications and any other host-level software, and reconfigure the replacement PM to match your original settings.

Note: When the replacement PM exits maintenance mode, the system automatically disables the PXE server on the primary node that was used for the replacement process.

To remove and replace a failed PM (with USB installation)

Use the following procedure to replace a failed PM and reinstall the system software by using a USB medium.

1. In the ztC Console, click **Physical Machines** in the left-hand navigation panel.
2. Select the appropriate PM (node0 or node1) and then click **Work On**, which changes the PM's **Overall State** to **Maintenance Mode** and the **Activity** state to **running (in Maintenance)**.
3. After the PM displays **running (in Maintenance)**, click **Recover**.

4. When prompted to select the type of repair, click **USB PM Replace - Initialize All Disks**.

   Caution: Selecting **USB PM Replace - Initialize All Disks** deletes all data on the replacement PM.

5. Click **Continue** to begin the replacement process. The system shuts down the PM in preparation for the system software reinstallation.
6. After the PM is powered off, install the replacement PM, if applicable:
   a. Disconnect and remove the old PM, and then install the replacement PM. Connect a monitor and keyboard.
   b. Reconnect the network cables to their original ports.
   c. Connect the bootable USB medium to the replacement PM, and then reconnect the power cable. If the PM does not automatically power on, press the power button.
7. As the replacement PM powers on, enter the firmware (UEFI) setup utility. In the **Save & Exit** menu, under **Boot Override**, select the **UEFI** entry for the USB medium to boot from the device one time during the next boot sequence. The PM restarts.

   Note: Use the **Boot Override** property to temporarily change the boot device instead of modifying the persistent **BOOT ORDER Priorities** in the **Boot** menu. The top boot priority must remain **UEFI Network** (default) to support the automated node replacement that is typically performed on ztC Edge systems.

8. Monitor the installation process at the physical console of the replacement PM.
9. At the **Welcome** screen, use the arrow keys to select the country keyboard map for the installation.
10. At the **Install or Recovery** screen, select **Replace PM, Join system: Initialize Data** and press **Enter**. The replacement process continues with no user interaction.

    Caution: Selecting **Replace PM, Join system: Initialize data** deletes all data on the replacement PM.

11. When the software installation is complete, the replacement PM reboots from the newly installed system software.

> **Note**: After the system software installation, the replacement PM may take up to 20 minutes to join the system and appear in the ztC Console.

12. As the replacement PM joins the system, you can view its activity on the **Physical Machines** page of the ztC Console. The **Activity** column displays the PM as **(in Maintenance)**, and then as **running** after the replacement is complete. The PM automatically exits maintenance mode and begins load balancing the VMs on the system.
13. If applicable, manually reinstall applications and any other host-level software, and reconfigure the replacement PM to match your original settings.

Related Topics

Maintenance Mode

Maintaining Physical Machines

The ztC Console

Physical Machines and Virtual Machines

The Physical Machines Page

Part 2: Supporting Documents

See the following support documents for release information, and reference and troubleshooting information.

- ztC Edge Release 1.3.0.0 Release Notes
- System Reference Information
- SNMP

8

Chapter 8: ztC Edge Release 1.3.0.0 Release Notes

These Release Notes are for ztC Edge Release 1.3.0.0 (updated at 10:19 PM on 4/29/2021). See the following sections:

- Accessing Stratus Knowledge Base Articles
- Bug Fixes
- New Features and Enhancements
- Important Considerations
- Known Issues
- Getting Help

Accessing Stratus Knowledge Base Articles

The **Stratus Customer Service Portal** provides a searchable **Knowledge Base** with technical articles about all Stratus products, including ztC Edge. In some cases, the Release Notes directly reference these Knowledge Base articles (for example, KB-*nnnn*). You can access the Customer Service Portal and the Knowledge Base articles by using your existing service portal credentials, or by creating a new user account, as follows.

To access the Knowledge Base

1. Log on to the **Stratus Customer Service Portal** at https://support.stratus.com.

   If needed, create a new account as follows:

   a. Click **Register Account**.

   b. Enter your company email address and contact info, and click **Register**.

Your company email address must include a domain name (for example, stratus.com) for a company that is a registered customer of Stratus.

c. 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 Stratus service representative.

2. In the service portal, click **Knowledge Base** in the left pane.
3. In the **Keyword Search** box, enter keywords associated with the information you need, and then click **Search**.

   To search for an article by its KB-*nnnn* number, click **Advanced Search**. Next to **Search by ID**, type the article ID number (*nnnn*) and click **Display**.

Bug Fixes

Fixed in ztC Edge Release 1.3.0.0

EV-48289,EV-48275,EV-47731,EV-48317,EV-48399: Fixed restart issues with an internal process.

Fixed in ztC Edge Release 1.2.1.0

EV-46765: Nodes may fail to return to service after power loss.

EV-47788: The eAC may become inaccessible after system power loss.

EV-47803: A spine restart may cause a false "Node has Rebooted" alert.

EV-47807: Refine the install checks that identify the ztC Edge 100i platform.

EV-47824: The fatal alert "A sensor on node0 (or 1) is indicating a fatal problem" can be seen when it should not be.

EV-47858: Insulate volume replication status and epoch values against node power loss.

EV-47865: The heartbeat configuration file /etc/ha.d/ha.cf can be truncated on node power loss resulting in heartbeat and spine failing to start and the node failing to join the cluster.

EV-47944: /etc/lvm/lvm.conf can be truncated after ztC node power loss.

New Features and Enhancements

Major new features and enhancements are listed below under the release in which they became available.

New in ztC Edge Release 1.3.0.0

ztC Edge Release 1.3.0.0 includes support for the following new features:

- Quorum Server (see Quorum Servers)
- Service heartbeat, for periodically reporting system health to your authorized Stratus service representative (see Configuring Remote Support Settings)
- REST API, for monitoring system health and statistics (see REST API)

New in ztC Edge Release 1.2.1.0

Release 1.2.1.0 is for bug fixes. See Bug Fixes for information.

New in ztC Edge Release 1.2.0.5

ztC Edge Release 1.2.0.5 includes the following new features:

- CentOS 7.5 host operating system for Stratus Redundant Linux
- CentOS 7.4 and 7.5 guest operating systems for virtual machines (see Compatible Guest Operating Systems)
- Selecting a keyboard map during or after deployment (see Deploying the System and Mapping Your Keyboard)

Also, beginning in Release 1.2.0.5, the online Help is available in German, Japanese, and Chinese.

Important Considerations

Upgrading to ztC Edge Release 1.3.0.0

To upgrade to the current release of ztC Edge, follow the instructions in Upgrading Stratus Redundant Linux Software Using an Upgrade Kit.

Caution: All PMs and VMs must be in good health before you upgrade the Stratus Redundant Linux software. Before starting an upgrade, examine the ztC Console to verify that there are no alerts indicating PM or VM problems.

Note: Eject any VCDs or USB media from the VMs before upgrading. If VCD or USB media is still connected to the VMs, it prevents the system from migrating the VMs and putting the PMs into maintenance mode for the upgrade process.

To determine the version of Stratus Redundant Linux running on a ztC Edge system, log on to the ztC Console for the system and check the system information in the masthead:

```
ocean.abc.com
IP: 123.109.50.34 | Asset ID: ze-12345
Version: n.n.n-nnn
```

If the software release is lower than Release 1.3.0.0, download the Stratus Redundant Linux 1.3.0.0 Upgrade Kit from the **Downloads** page at https://www.stratus.com/services-support/downloads/?tab=ztcedge and upgrade the software on the system as described in Upgrading Stratus Redundant Linux Software Using an Upgrade Kit.

Using Intel Active Management Technology (AMT) for Lights-Out Support

ztC Edge systems incorporate Intel Active Management Technology (AMT) lights-out support for remote power management, remote console, and remote media. For important information about AMT configuration and restrictions, see KB-8219.

Redeploying a ztC Edge System

Redeploying a ztC Edge system allows you to clear the network settings of the system to prepare it for deployment on a different network. If you need to move a system to a new network, redeploy and shut down the system **before** moving it, as described in Redeploying a ztC Edge System.

When you redeploy a system, note the following issues and workarounds:

- Disable any NFS/CIFS shares before redeploying a system.

  Active NFS/CIFS shares interfere with the redeployment feature. Disable the shares until you finish configuring the network settings on the new network.

- System reboot needed when setting new static system IP address.

  A system loses access to the secondary node if you redeploy and shut down the system, start it up in a new location, and then configure a new static system IP address. To regain access to the secondary node, reboot the system by opening the **System** page and clicking **Reboot**. Rebooting the system refreshes the gateway settings on the secondary node and allows it to connect to the system.

- If you already moved a system to a new network, but you forgot to redeploy it first, see KB-8283 for instructions on redeploying the system.

- If you need to redeploy a single, used node as the first node in a new system, or as the secondary node in a different system, contact your Stratus service representative and refer to KB-9391.

Deploying ztC Edge Nodes at Separate Physical Sites

When you deploy a ztC Edge system, you must deploy both nodes at the same site and directly connect the A-links between the blue (A2 •) and yellow (A1 • •) network ports of each node. If you want to set up a ztC Edge system in an Automated Local Site Recovery (ALSR) configuration, where each node is located at a separate physical site for increased redundancy, contact your authorized Stratus service representative for assistance. Because of the geographic separation, an ALSR configuration requires careful planning of component placement and networking topologies.

Other Important Considerations

For other important considerations about ztC Edge systems, see Important Physical Machine and Virtual Machine Considerations.

Known Issues

Do Not Click the Remove Button on the Physical Machines Page

Do not click the **Remove** button on the **Physical Machines** page of the ztC Console. If you remove a node from the system, you will be unable to restore the node to service and also unable to create new VMs. If you have already clicked the **Remove** button, contact your Stratus service representative to restore the node to service and correct the VM issue.

NFS or SMB/CIFS Mounts May Delay Failover

If an NFS or SMB/CIFS share is mounted on a ztC Edge system and the ibiz0 (P1) network connection is lost, you may lose access to the ztC Console for several minutes, and you may need to refresh the web browser to restore the connection. To avoid failover delays, disable any shares in the host operating system when they are not in use.

Importing or Exporting a VM Through a Windows Share (CIFS/SMB) May Fail

When using a Windows share to import or export VMs, the share may not show up as a valid option or may not allow connection because the operating system hosting the share does not support SMB v1.0. For more information about this issue, see KB-9541.

VMs Reported as Broken Instead of Degraded When A-Link Is Offline

If an A-link cable or network is disconnected on one node, the state of a VM on that node may be reported as broken () in the ztC Console, even though the VM still has another active A-link connection. The availability

of the VM is unaffected.

Ejected VCD Still Displayed in a Linux-based VM Console

If you use the ztC Console to eject a VCD from a VM running a Linux-based guest operating system, the VCD may still be displayed in the guest operating system. If needed, you can eject the VCD in the guest operating system to stop displaying the VCD.

Cannot Re-mount USB Media After Ejecting from Windows File Explorer in a VM

If you eject USB media from File Explorer in a Windows VM, you cannot re-mount the media without first physically disconnecting the media from the ztC Edge system.

There are two methods for ejecting media from a Windows-based guest operating system:

- Clicking **Eject** in File Explorer

  If you eject media from File Explorer, you must eject the media in the ztC Console and then physically disconnect the media from the ztC Edge system before reconnecting it.

- Clicking **Safely Remove Hardware and Eject Media** in the taskbar

  If you eject media from the taskbar, you must eject the media in the ztC Console, but you do not need to physically disconnect the media from the ztC Edge system before reconnecting it.

VM Network Remains Offline After Using P2V/V2V to Migrate from a VMware Source

If you use the P2V client (virt-p2v) to migrate a PM or VM from a source that is running VMware, the network driver might fail to install properly during the migration process, which leaves the VM network offline. If this happens, manually install the network driver in Device Manager. For instructions, see the **Troubleshooting** section of Migrating a Physical Machine or Virtual Machine to a System.

VM Network Remains Offline After Importing a Windows-based VM

If you import a Windows-based VM with the **Import/Restore Virtual Machine Wizard**, the VirtIO network driver may fail to install properly the first time you start the VM, which leaves the VM network offline. If this happens, manually install the network driver in Device Manager. For instructions, see the **Troubleshooting** section of Importing an OVF File.

Ensure that Both Nodes Are Connected to the Network Before Creating a VM

Before creating a VM, always ensure that the network you assign to the VM has an active network connection in each node of the system. You can verify the state of the networks on the **Networks** page of the ztC

Console.

If you add a VM while one of the paired network interfaces is offline, the VM will not start. To start the VM, connect an active network cable to the offline interface, wait a few seconds, and start the VM.

Ensure that Both Nodes Are Running Before Removing a VM

To avoid downtime, before removing a VM, always ensure that both nodes are in the **running** state and that neither node is in maintenance mode or in the process of synchronizing.

If you remove a VM while one node is shut down in maintenance mode, the operation may stall with the error, "There was a problem deleting *vm-name*." To correct the problem, you must restart the only running node.

Invalid VM Names Prevent Some VM Operations from Executing

The ztC Console may improperly validate VM names to ensure that they meet the software requirements. If you are creating, reprovisioning, or renaming a VM with an unsuitable name, the operation may fail to execute.

To avoid any problems with VM operations, ensure that your VM names meet the following requirements:

- A VM name must start with a character within the range [a-z], [A-Z], or [0-9].
- A VM name cannot start with the prefixes `Zombie-` or `migrating-`, which are reserved strings.
- A VM name can contain characters only within the range [a-z], [A-Z], or [0-9]; or _ (underline), - (minus), or . (dot).
- A VM name must be 1 to 64 characters in length.

Some Browsers Unable to Connect a VNC When Using https

If you are connected to the ztC Console using an **https** URL in a Microsoft Internet Explorer or Mozilla® FireFox® browser, and you click **Console** after selecting a running VM from the **Virtual Machines** page, the message **VNC: Unable to connect, retrying in *n* seconds** may appear. To enable the VNC connection, click the https link to the VNC console page in the upper right-hand corner of the masthead, and continue with the appropriate procedure below:

- In Internet Explorer, the **Security Alert** wizard appears:
    a. Click **Continue to this website (not recommended)**.
    b. Click **OK**.
- In FireFox, the **Your connection is not secure** window appears:

a. Click **Advanced**. A message about an invalid security certificate appears.

b. Click **Add Exception**. The **Add Security Exception** dialog box appears with the console's location in **Location**.

c. Click **Confirm Security Exception**.

The VNC console appears.

Cannot Mount USB Media with exFAT File System

The Stratus Redundant Linux software on ztC Edge systems does not support the exFAT File system. If you need to mount a USB medium to import or export a VM, as described in Mounting a USB Device or Network-mounted Folder on the ztC Edge System, format the device with NTFS. (By default, most USB media are formatted with the FAT file system, which has a limited file size of 4 GB that may be too small for most VMs.)

Reboot Required when Changing Node IP Address or Netmask Network Settings

When you change the IP address or netmask settings of an ztC Edge node as described in Configuring IP Settings, you must reboot the node or the system for the changes to take effect. If you do not reboot the node or the system, both the new and old settings will remain in effect, which may cause routing or connection issues.

Changing the MTU of Business Networks Causes VM Migration and Possibly Node Failover

If you change the maximum transmission unit (MTU) setting of either `network0` (P1) or `network1` (P2), the system automatically migrates the VMs from one node to the other. If you change the MTU for `network0` (P1) specifically, the system also automatically fails over from the primary node to the secondary node.

To avoid this issue, avoid changing the MTU setting of the business networks, or change the MTU only during a planned maintenance period.

Getting Help

If you have a technical question about ztC Edge systems, you can find the latest technical information and online documentation at the **Downloads** page at https://www.stratus.com/services-support/downloads/?tab=ztcedge. You can also search the **Knowledge Base** in the **Stratus Customer Service Portal** at https://support.stratus.com.

If you cannot resolve your questions with these online resources, and the system is covered by a service agreement, contact your authorized Stratus service representative. For information, see the **ztC Edge Support** page at https://www.stratus.com/services-support/customer-support/?tab=ztcedge.

9

Chapter 9: System Reference Information

See the following topics for reference information

- Compatible Guest Operating Systems
- Important Physical Machine and Virtual Machine Considerations
- Accessing Knowledge Base Articles
- Creating an ALSR Configuration
- REST API

Compatible Guest Operating Systems

The following are compatible as guest operating systems for virtual machines (VMs) running on ztC Edge systems.

Operating System	Version	Boot Firmware Interface
CentOS 7	7.4, 7.5 (all 64-bit)	BIOS
Microsoft Windows Server 2019 (Standard, 2 Core Datacenter)	64-bit	BIOS
Microsoft Windows Server 2016 (Essentials, Standard, Datacenter)	64-bit	BIOS

Operating System	Version	Boot Firmware Interface
Microsoft Windows Server 2012 (Foundation, Essentials, Standard, Datacenter)	64-bit, 64-bit R2	BIOS
Microsoft Windows Small Business Server 2011 (Standard, Essential, Premium Add-On)	64-bit	BIOS
Microsoft Windows Server 2008 (Web, Small Business, Standard, Enterprise, Datacenter)	32-bit, 64-bit R2 only	BIOS
Microsoft Windows 10 Desktop	64-bit	BIOS
Microsoft Windows 8.1 Desktop (Enterprise)	64-bit	BIOS
Microsoft Windows 8 Desktop (Enterprise)	64-bit	BIOS
Microsoft Windows 7 Desktop	32-bit, 64-bit	BIOS
Ubuntu	12.04 LTS, 14.04 LTS, 16.04 LTS (all 64-bit) 17.10.1 Server 64-bit 18.04 Server 64-bit	BIOS

Important Physical Machine and Virtual Machine Considerations

For optimal implementation of physical machines and virtual machines, be aware of the configuration maximums and requirements described in the following sections:

- Virtual Machine Recommendations and Limits
- Important Considerations

Virtual Machine Recommendations and Limits

Virtual machines (VMs) require certain CPU core resources.

Recommended Number of CPU Cores

Stratus recommends using only as many threads for workloads as physical threads on a ztC Edge system. The ztC Edge 100i system has 8 physical threads, total.

Examples

Four 2-vCPU HA guests typically require 8 threads, total.

Two 3-vCPU HA guest and one 2-VCPU HA guest typically require 8 threads, total.

Two 4-vCPU HA guests typically require 8 threads, total.

One 8-vCPU HA guest typically requires 8 threads, total.

Important Considerations

Note the following important considerations.

Feature	Comment
USB Devices	USB keyboards, CD/DVD drives, disk drives, and thumb drives are supported for importing/exporting VMs and for system restoration.
Console Connectivity	Each PM's text console is available for a CentOS operating system. However, VGA mode is not supported; that is, the PM must be at run-level 3 and cannot run at run-level 5. See "System Management" below.
System Management	ztC Edge system management **does not work** at run-level 5.
Volumes	For exporting, importing, or restoring a volume, the maximum volume size is 2TB.

Creating an ALSR Configuration

This topic and its subtopics describe how to create an automated local site recovery (ALSR) configuration. For general information about quorum servers, see Quorum Servers as well as ALSR and Quorum Service.

> **Note**: **Before** you create an ALSR configuration, read this topic and all of its subtopics and then plan your ALSR configuration, as described in the topics. Create the configuration only after you are certain that your planned configuration complies with the information in this topic and its sub-topics.

An ALSR configuration exists if either of the following is true:

- The two nodes of a dual-node system are connected using network infrastructure rather than direct cables.
- The length of the A-Link (direct connect) cables connecting the two nodes is greater than 10m (for example, in two separate buildings within a campus).

These configurations provide better disaster tolerance and hardware redundancy as well as redundancy of physical computer rooms and the buildings containing them.

Stratus recommends that an ALSR configuration include a third computer, which is a quorum server. The quorum server is located in a physical location that is removed from the physical location of both node0 and node1.

> **Note**: This topic and its subtopics describe an ALSR configuration with a quorum server. Stratus highly recommends that an ALSR configuration include a quorum server. If you want to consider creating an ALSR configuration without a quorum server, access the Knowledge Base to search for the article *Considerations if deploying ALSR without quorum* (KB-9682), and also contact your authorized Stratus service representative. For information about accessing Knowledge Base articles, see Accessing Knowledge Base Articles.

Because of the geographic separation of these physical machines, creating an ALSR configuration requires careful planning of component placement and more complex networking topologies.

The topics below describe how to create a ALSR configuration. To perform the procedures in the topics, you should be familiar with ztC Edge software and the hardware it runs on, and you should be familiar with the network infrastructure of your system and its location.

> **Note**: These topics cannot describe every vendor and model of network switches, routers, and other hardware. Consult the documentation that pertains to your infrastructure if you need more information about how to configure it according to the requirements in these Help topics.

- Creating the Configuration
- Meeting Network Requirements
- Locating and Creating the Quorum Server
- Completing the Configuration
- Understanding Quorum's Effect on System Behavior

The following table lists and defines terms associated with creating an ALSR configuration.

Term	Meaning
Active node	The node where a guest VM is currently running. Each guest VM may have a different active node. The opposite of *active* is standby (see Standby node).
A-Link	Availability link. A direct network connection between the two computers that form a ztC Edge system. (The computers of a system are also referred to as *physical machines* (PMs) or *nodes*.) A-Links must be connected point-to-point, and A-Link traffic cannot be routed. A ztC Edge system requires two A-Links. On some systems, these connections have blue and yellow cables (and ports). You can use VLAN connections for A-Links in a distributed local site deployment (see VLAN).
Alternate quorum server	The alternate quorum server is used when the preferred quorum server is not available (see Preferred quorum server).
Automated local site recovery (ALSR)	An ALSR configuration exists if either of the following is true: • The two nodes of the ztC Edge system are connected using network infrastructure rather than direct cables. • The length of the A-Link (direct connect) cables connecting the two nodes is greater than 10m (for example, in two separate buildings within a campus). An ALSR configuration is typically used to provide better disaster tolerance, at the expense of more network setup and more extensive configuration options. An ALSR configuration requires a third computer, which

	is a quorum server (see Quorum server).
AX	The container layer that resides within the ztC Edge system and controls the behavior of the guest VM. AX is responsible for keeping a VM synchronized between the active node and the standby node. Each VM has its own AX pair (see VM, Active node, and Standby node)
Business network (ibiz)	A network connection from the ztC Edge system to a LAN that also has other traffic that can include management messages as well as traffic for applications and other clients and servers. The ztC Edge system typically has two ports for business network connections. Business networks can be assigned to one or more guest VMs for their use, or to no guest VMs. You must connect the first business network (ibiz0) to a LAN so that you can manage the system from a web browser.
Fault	Any potential degradation in a system's ability to execute a guest VM (see VM). Disk failure, network loss, or power outage are all examples of faults detected by the system.
Node0 and node1	The two computers that form the ztC Edge system are labeled internally as node0 and node1. (These computers are also sometimes referred to as physical machines or PMs.) The choice of node0 and node1 is arbitrary and is made when the system is configured for the first time. Constant traffic flowing between node0 and node1 communicates state information for the system as well as for each running guest VM (see VM).
Preferred quorum server	The preferred quorum server is used when it (the preferred quorum server) is available. If the preferred quorum server is not available, the alternate quorum server (if it exists) is used (see Alternate quorum server).
Primary node	When the system's computers are paired, only one computer responds to management messages. This computer is the primary node. The System IP address, which is assigned when the system is initially deployed, applies to the primary node. The primary node can switch between node0 and node1 as different fault conditions occur (see Fault). Note that the

	primary node is not necessarily the active node for a guest VM (see Active node and VM).
priv0	A network for private management traffic between the two nodes. For more information, see A-Link and Private Networks.
Quorum server	A third computer that helps arbitrate which AX should be active for each guest VM (see Active node and VM). Correct use of a quorum server is the only guaranteed way to prevent split-brain conditions (see Split-brain).
RTT	Round-trip time. The elapsed time required for a network message to travel from a starting point to a destination and back again. The time is typically measured in milliseconds (ms).
Split-brain	The condition that occurs when both AX's of a guest VM's AX pair are active simultaneously, which produces divergent copies of data within each active guest (see AX and VM). Split-brain can occur when all communication paths between node0 and node1 are disconnected (see Node0 and node1). Using the quorum service prevents a split-brain condition (see Quorum server).
Standby node	The node that is not the active node for a guest VM. The standby node is kept synchronized through AX communications via A-Link connections (see AX and A-Link). The AX pair for each guest VM determines which node is active and which is standby (see Active node).
System management	System management is the layer within Stratus Redundant Linux software that is responsible for maintaining the overall state of the system. Determining which node is primary is part of system management (see Primary node). System management is also responsible for displaying information within the ztC Console.
UPS	Uninterruptable power supply. An external battery backup for electrical equipment that prevents short power outages from affecting availability.

VLAN	Virtual LAN. A VLAN is a set of devices on one or more LANs that are configured to communicate as if they were attached to the same cabled network, when in fact they are located on different LAN segments. VLANs are configured at the network infrastructure level, not within the ztC Edge system. In anAutomated local site recovery (ALSR) configuration, the A-Link connections are implemented as isolated VLANs (see A-Link).
VM	Virtual Machine (also referred to as a guest). A system typically has one or more VMs (or guests) allocated and running applications via guest operating systems.

Creating the Configuration

To create an ALSR configuration, first consider the configuration of a typical ztC Edge system configuration and the VLAN requirements of an ALSR configuration. Then, observe a well-planned ALSR configuration, which includes a quorum server, and become familiar with the configuration's VLAN requirements. You must also become familiar with the entire process of deploying a typical ztC Edge system and then creating an ALSR configuration. The sections below provide this information.

A Typical ztC Edge System

In a typical ztC Edge system configuration, two PMs are directly connected by a pair of network cables for A-Links. One A-Link typically serves as the private network (priv0). The two PMs have additional network connections for business networks, which the ztC Console and guest VMs hosted by the system use. The following figure illustrates a typical configuration.

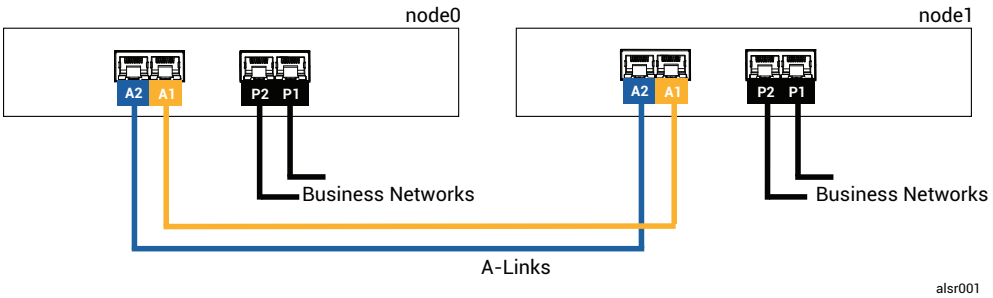

The physical distance between the PMs in a typical configuration is limited by the length of a single A-Link network cable, which is approximately 33 ft (10m). This distance may be significantly shorter when the physical environment and ambient electrical noise is accounted for.

An ALSR Configuration With a Quorum Server

A well-planned ALSR configuration consists of the two nodes at two different locations, and a third computer that runs the quorum service at a third location. All of these computers are networked together with appropriate network switching equipment, so that no single point of failure exists within the ALSR configuration. The following figure illustrates such a configuration, which includes node0 at Site A, node1 at Site B, and the quorum server at Site C.

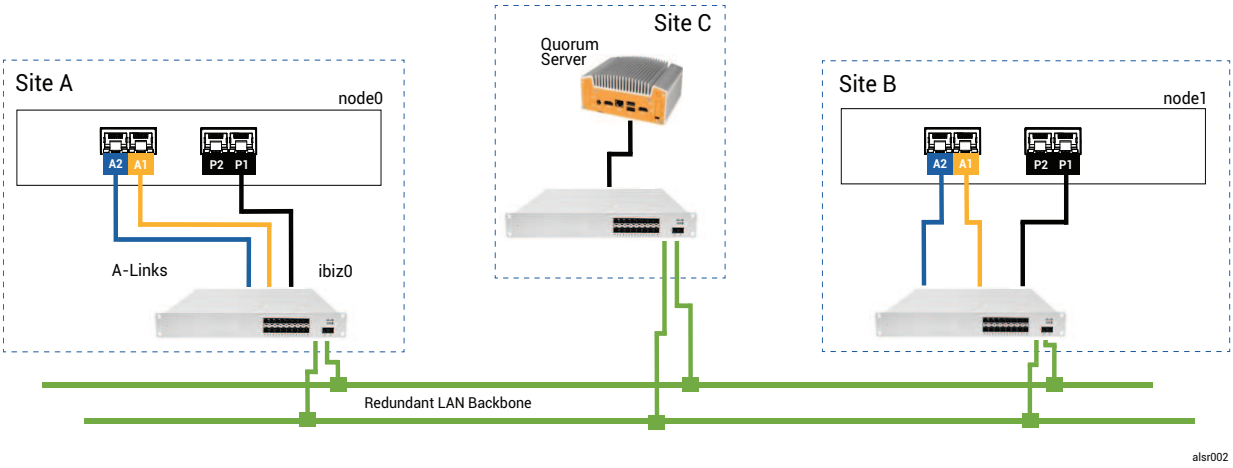

Notes:

1. Each A-Link should be connected on its own VLAN configured between switch A and switch B.
2. DNS servers and gateways are not included in the illustrations, for clarity, but you must ensure that the ALSR configuration includes a connection to a DNS server and a gateway in the event of a network failure.
3. For maximum protection, you should install redundant switches at each at site though the illustration does not show these switches. For the illustrated configuration, site A and site B would *each* include two switches. The A-Links are routed through one switch and the business networks are routed through the other switch. If possible, use different circuits to power the switches or use a UPS to prevent brief power loss failures.

ALSR VLAN Requirements

The A-Link connections between switch A and switch B require a VLAN configuration on the switches. A-Link traffic is not routable, and the connection should emulate a single long network cable. Each A-Link must be isolated on its own VLAN.

If you cannot create VLANs between the switching equipment, you can use Ethernet-to-fiber media converters to create a longer fiber connection between the two PMs. However, you should not route the two A-Link fiber connections through the same physical conduit, as this creates a single point of failure.

Additionally, the quorum service computer should not share a switch with either node0 or node1 because a shared switch creates a single point of failure.

See Meeting Network Requirements for more information about the latency requirements of the A-Links and quorum connections.

From Initial Deployment to Completing the ALSR Configuration

When creating an ALSR configuration, you should first deploy and register a typical ztC Edge system, initially without the ALSR configuration. The figure in A Typical ztC Edge System illustrates this system. For simplicity, install the nodes side-by-side, using the provided cables. See Getting Started.

After the typical system is operating normally, create the ALSR the configuration.

1. Read Creating an ALSR Configuration and all of its subtopics, if you have not already done so.
2. Install the quorum computer and enable the quorum server. Comply with all information in:
   - An ALSR Configuration With a Quorum Server
   - ALSR VLAN Requirements
   - Meeting Network Requirements
   - Completing the Configuration
3. Verify that the quorum server has access to both nodes.
4. Properly shutdown one node. See Shutting Down a Physical Machine.
5. Relocate the shutdown node to the far site.
6. Connect the infrastructure. The ALSR-configuration illustration above shows the connections, which include:
   - The priv0 connection to port **A2**
   - The second A-Link connection to port **A1**
   - The ibiz0 connection to port **P1**
7. Power on and (re-)join the nodes. See Powering On a Physical Machine.
8. Verify the configuration. Ensure that:
   - The shared networks pair properly—In the ztC Console, navigate to the **Networks** page and ensure that the state of each network is green-checked. If necessary, troubleshoot any infrastructure problems.
   - Quorum connections are remade—In the console, navigate to the **Quorum Servers** page by clicking **Preferences** and then **Quorum Servers**. Ensure that the state of the quorum server is green-checked. If necessary, troubleshoot any infrastructure problems.
   - The primary node can shift from node0 to node1, and the console can connect in both configurations—Place each node in Maintenance Mode (see Maintenance Mode).
9. (Re-)join the VMs—Migrate the VMs from node to node (see Migrating a Physical Machine or Virtual Machine to a System). Verify the correct network failover of VM networking.
10. Assess the status of network and validate Ethernet failover (see The Networks Page).

Meeting Network Requirements

This topic describes the network requirements and considerations of A-Links, business networks, the quorum server connections, and the management network for a successful ALSR configuration .(For general information about these networks, see Network Architecture.)

Prerequisite: Plan and create an ALSR configuration by first reading Creating an ALSR Configuration and following its instructions, if you have not already done so.

A-Link network connections must meet the following requirements:

- The A-Links use IPv6 addressing.
- Each A-Link must be connected on its own VLAN. A-Link traffic is not routable.
    - HA VMs require less than 10ms RTT A-Link latency (available on all ztC Edge systems).
    - You need to provide enough bandwidth to meet the needs of all VMs on the system, and you need to provide a speed of at least 1Gb per A-Link.
    - When planning your network infrastructure, you need to account for the uplink bandwidth between the switch and the network backbone across all the ports in use on that switch.

  If these requirements are not met, guest VMs may run more slowly due to limited synchronization bandwidth between the two nodes.

The first business network (ibiz0) is used for communication between the nodes and to the quorum server. The ibiz0 network must meet the following requirements:

- The two nodes must be on the same subnet.
- The network must allow IPv6 multicast traffic between the two nodes.
- The two nodes can access the quorum server using IPv4 network addressing.

Network connections for the quorum server must meet the following requirements:

- Access to the quorum service must be provided using ibiz0, using IPv4 network addressing.
- Two UDP ports must be open and available for communication between the nodes and the quorum service, including in the firewalls. By default, these ports are 4557 and 4558. If you want to change these ports, see Configuring the Quorum Service Port (on the quorum computer) and Configuring the Quorum Server Within the ztC Console.
- Latency between a ztC Edge node and the quorum computer should be less than 500ms RTT.

- Throughput is not an important consideration. 10Mb Ethernet, or even T1 bandwidth is adequate.
- Quorum computers are common to all VMs on the same ztC Edge system.
- Quorum computers may be shared among many ztC Edge systems.
- Quorum computers must never be implemented as a VM on the same ztC Edge system that uses it.
- Use different network infrastructure, don't share. A ztC Edge node should not depend on a gateway or switch/router on the partner node site for sustained access to a quorum services computer.

Note: Do not implement the quorum service as a guest VM on a different pair of nodes; a failure on those nodes would cause the VM running the quorum service to failover, which would create unnecessary complications for network topology and fault management. Additionally, a second quorum computer is needed to manage quorum for the ztC Edge system that is running the quorum service. .

Management network connections must meet the following requirements:

- By default, the management network is shared with a business network. In this case, all requirements for a business network apply.
- Configure gateways to a business LAN for remote management.

Locating and Creating the Quorum Server

In a well-planned ALSR configuration, a third computer hosts the quorum service. The quorum service processing requirement is small, so any other existing computer or VM that meets all network and operating requirements can host the quorum service. An effective quorum server depends upon the location of the quorum computer within your network.

After you have determined an effective location for the quorum computer (and an alternate quorum computer, if desired) and ensured that the computer meets the requirements of the quorum service, you can create the quorum server.

Prerequisite: Plan and create an ALSR configuration by first reading Creating an ALSR Configuration and following its instructions, if you have not already done so.

Locating the Quorum Computer

Locate the first quorum computer in a third site within your network, as An ALSR Configuration With a Quorum Server illustrates. If a third site is not available, locate the quorum computer in a physical location that is different from the physical location of node0 and node1. Locating the quorum computer in a unique site maximizes the chance of the system surviving a problem that causes the loss of both nodes and the quorum computer (for example, a transient electrical, plumbing, or other problem that causes loss of network connectivity).

You should connect the quorum computer to an electrical circuit that is different from the electrical circuit that powers node0 or node1. In addition, you should connect the quorum computer to a UPS unit.

Caution: If both AX's lose connectivity with the quorum server, they will attempt to select an alternate quorum server. If no quorum server can be selected, the VM is downgraded to simplex mode, to prevent a split-brain condition if another failure occurs.

If one node shuts down and the VM (AX) on the remaining node cannot reach either the quorum server or its peer, it will shut itself down to avoid the risk of a split-brain condition.

When locating the quorum computer:

- Ensure that the quorum computer does not share a switch (or router) with either node0 and node1.
- Do **not** use a guest VM within the ztC Edge system to run the quorum service.

See Understanding Quorum's Effect on System Behavior for a description of system behavior and failure modes.

Adding an Alternate Quorum Computer

You can add another quorum computer (with its switch) to your system to create an alternate quorum service. The most common use of an alternate quorum server is when, for example, operating system updates are being applied to the preferred quorum computer. When the preferred quorum computer restarts, the alternate quorum computer is selected and no downgrade occurs. When the preferred quorum is recovered, the selection moves back to the original preferred quorum computer.

When creating a second quorum service, you must follow all of the requirements for the network and quorum placement. If both nodes can communicate with each other and with the same quorum server (either the preferred or alternate quorum server), the system can maintain VM redundancy, even if one quorum connection is

lost. Preferred quorum server selection occurs when both nodes have access to each other and to the preferred quorum server. Thus, if the preferred quorum service is lost at the same time a node is lost, the remaining node shuts down the VM even if a second, non-preferred quorum service is available. However, if the preferred quorum service is lost *before* a node is lost, and if both nodes can continue to contact the alternate quorum server, the selection is moved to the alternate quorum server. Fault handling occurs in a context of the selected quorum server only.

If you create an alternate quorum service, you need to add a second quorum IP address when adding the quorum service in the ztC Console.

Quorum Computer Requirements

You can install quorum service software on any general-purpose computer, laptop, or VM that is running the Windows operating system and that meets these requirements:

- The computer can continually remain powered on and connected to the network such that the ibiz0 network of the ztC Edge system can always access the quorum server.
- The computer has a static IPv4 network address. Do not use DHCP.
- The operating system is Windows Server 2016, Windows Server 2012, Windows Server 2008, Windows 7, or Windows 10; Embedded versions of the Windows OS are not supported.
- A minimum of 100 MB disk space is available.
- Two UDP ports must be open and available for communication between the nodes and the quorum service, including in the firewalls. By default, these ports are 4557 and 4558. To change these ports, see Configuring the Quorum Service Port (on the quorum computer) and Configuring the Quorum Server Within the ztC Console.

Downloading and Installing the Quorum Service Software

After you have determined an appropriate location for the quorum computer, download and install the required software to create the quorum server.

To download and install the quorum server software

1. Open the **Downloads** page at https://www.stratus.com/services-support/downloads/?tab=ztcedge.
2. Scroll down to the **Drivers and Tools** section and then click **Quorum Service** to download the quorum server software installer file to the quorum server.

3. On the quorum server, double click the installer file.
4. Move the downloaded file to an accessible location.
5. Log in to the quorum computer.
6. Navigate to the quorum service installer and double-click it.
7. Follow the prompts to complete the installation.

The product name *everRun* may appear when installing the quorum service.

Note: When upgrading to a more recent version of quorum server software, you do **not** need to uninstall the previous version.

Completing the Configuration

After you have created the ALSR configuration, change the quorum service port, if necessary. Then, enable quorum within the ztC Console. Finally, verify the configuration and (re-)join VMs.

Prerequisite: Plan and create an ALSR configuration by first reading Creating an ALSR Configuration and following its instructions, if you have not already done so.

Note: The port configured for quorum service on the quorum computer and the port configured for the quorum server within the ztC Console must be the same port numbers. If you change the quorum service ports on the quroum computer, you must change the quorum service ports on all ztC Edge systems (using the ztC Console) that connect to that quorum computer so that both the quorum computer and the ztC Edge systems use the same port numbers. See Configuring the Quorum Server Within the ztC Console.

Configuring the Quorum Service Port

By default, the quorum service listens on UDP port 4557.

In most cases, you do not need to change the default port. However, you can change the port, if the network configuration requires you to:

To change the port number on the quorum server

1. Log on to the quorum computer using an account with administrative privileges.
2. Open a command window in administrative mode.
3. Stop the quorum service by typing:
   ```
   net stop sraqserver
   ```
4. Change the port by typing (replacing *nnnn* with the new port number):
   ```
   sraqserver -install nnnn
   ```
5. Restart the quorum service by typing:
   ```
   net start sraqserver
   ```

Verifying the Quorum Service Port

If you need to verify the quorum service port, check this Windows registry key:

```
HKEY_LOCAL_
MACHINE\SYSTEM\CurrentControlSet\Services\SraQserver\Parameters\
QSServerPortForReceive
```

Configuring the Quorum Server Within the ztC Console

Once the quorum service is running, you should enable the quorum service within the ztC Console. You can also remove a quorum server.

To enable the quorum service:

1. Login to the ztC Console with an account that has administrative privileges.
2. Click **Preferences** in the left-hand navigation panel, to open the **Preferences** page.
3. Click **Quorum Servers**. The quorum configuration page opens.
4. Click **Add Quorum Server** at the left side of the page.
5. In the **Add Preferred Quorum Server** dialog box, enter the following values (if a preferred quorum server already exists, the **Add Alternate Quorum Server** dialog box appears):
   - **DNS or IP Address**–Type the fully-qualified **DNS** host name or **IP address** for the preferred quorum server.
   - **Port** –The default port is 4557. Type a port number if you need a port that is different from the default. You need to type only one port number. The quorum service will open the port number

for **Port** and the next port (for example, 4557 and 4558)

> **Note**: The port number must match the port that the quorum service is listening on. (If necessary, you can change the port on the quorum server.)

Click **Save** to save the values.

6. Repeat steps 4 and 5 to configure a second, alternate quorum server. Stratus recommends configuring two quorum servers.
7. To enable quorum service, select the **Enabled** check box and click **Save**.

Changes to the quorum configuration do not effect running VMs. You must stop and restart any running VMs after changing the quorum configuration.

To remove a quorum server

> **Caution**: If you remove the preferred quorum server, the alternate quorum server becomes the preferred quorum server. If no alternate quorum server exists, removing the preferred quorum server automatically disables quorum service.

1. Navigate to the **Preferences** page of the ztC Console.
2. Click **Quorum Servers**.
3. Locate the entry for the quorum server you want to remove.
4. In the right-most column, click **Remove**.

> **Note**: If a VM is using the quorum server that you are removing, you must reboot the VM so that it no longer recognizes the quorum server, which allows the removal process to finish. The VM will downgrade to simplex mode until it is restarted with no quorum servers configured.

Verify the Configuration and (Re-)Join VMs

Verify the configuration and (re-)join VMs. Follow the appropriate steps in From Initial Deployment to Completing the ALSR Configuration.

Understanding Quorum's Effect on System Behavior

A quorum server in an ALSR system changes the system's availability and recovery behavior. To understand the quorum's effect on system behavior, you first need to understand the behavior of a system that does not have a quorum server.

Prerequisite: Plan and create an ALSR configuration by first reading Creating an ALSR Configuration and following its instructions, if you have not already done so.

A ztC Edge system is designed to provide high availability for one or more guest VMs, which allows the VMs to continue to run even during failures that would otherwise create application downtime. The ztC Edge system can continue to run guest VMs even with, for example, the loss of a single network connection, a hard disk, or even an entire computer.

However, if more catastrophic faults occur (for example, the loss of all possible network paths), the ztC Edge system attempts to determine the overall state of the total system. The system then takes the actions necessary to protect the integrity of the guest VMs.

The following examples illustrate the system's process during a catastrophic fault.

Example 1: A System Without a Quorum Server Experiences a Split-brain Condition

In this ALSR example, the ztC Edge system includes node0 and node1, but does not include a quorum server. Operation is normal; no faults are currently detected. The two nodes communicate their state and availability over the A-Link connections, as they do during normal (faultless) operation. The following illustration shows normal connections.

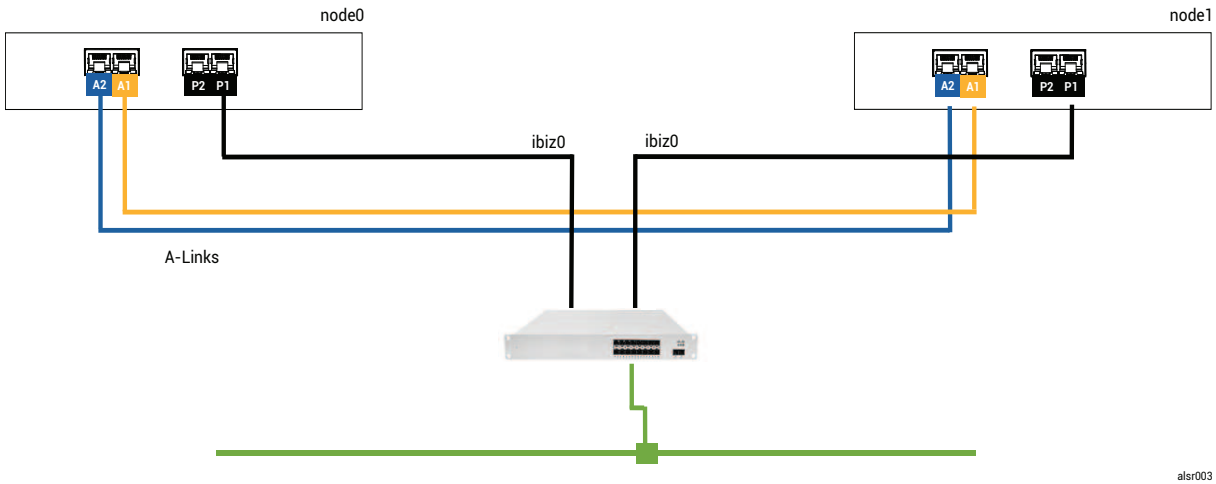

A Catastrophic Fault

A careless fork-truck operator crashes through the wall, severing all of the network connections (both business and A-Links), while leaving the power available and the system running. The following illustration shows the fault condition.

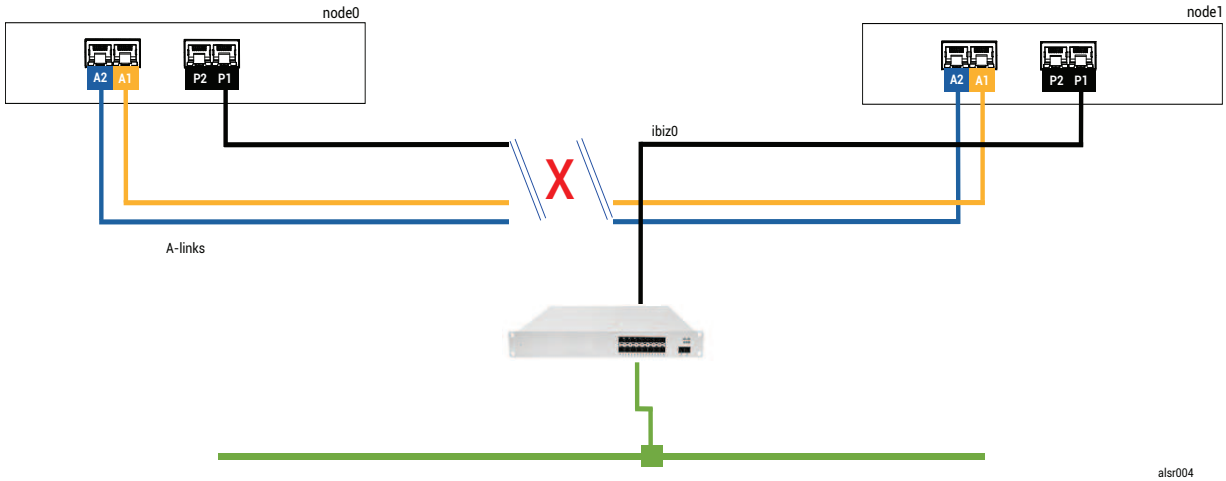

Fault Handling

The two nodes handle the fault, as follows:

- Node0—The AX on node0 detects the loss of both A-Links as well as all other network paths. Since the node0 AX can no longer detect the presence of its partner, the node0 AX becomes active and runs the guest VM. The application inside the guest VM continues to run, perhaps in a limited capacity due to the loss of the network.
- Node1—The AX on node1 also detects the loss of both A-Links, but ibiz0 remains available. As its partner does not respond to messages on ibiz0, the node1 AX is now active. The application inside the guest VM continues to run, perhaps not noticing any problems with the system.

From the perspective of an application client or an external observer, the guest VMs are both active and generate network messages with the same return address. Both guest VMs generate data and see different amounts of communication faults. The states of the guest VMs becomes more divergent over time.

Recovery and Repair

After some time, network connectivity is restored: the wall is repaired and the network cables are replaced.

When each AX of the AX pair realizes that its partner is back online, the AX pair with the fault handler rules choose the AX that continues as active. The choice is unpredictable and does not include any consideration for which node's performance was more accurate during the split-brain condition.

The data that was generated from the (now) Standby node is overwritten by the resynchronization of the Active node, and thus the data on the (now) Standby node is lost forever.

After a split-brain condition, the system requires several minutes to resynchronize, depending on how much disk activity needs to be sent to the standby node. If several guest VMs are running with different Active nodes, synchronization traffic may occur in both directions.

> **Note**: In some cases, the ztC Edge system may not be able to determine the best way to proceed after a catastrophic fault. In this case, a person needs to recover the system. The recommended recovery method is to use the ztC Console to shut down and reboot one node while the other node continues to run. This method typically forces the running node to become Primary and the AX on that node becomes Active. After the running node becomes Primary, a person can power on the other node. Do not shut down either node if resynchronization is already in progress.

Example 2: An ALSR System With a Quorum Server Avoids a Split-brain Condition

In this ALSR example, the ztC Edge system includes node0 and node1 with connections identical to those of the system in Example 1. In addition, the system in Example 2 includes a quorum server. The following illustration shows these connections.

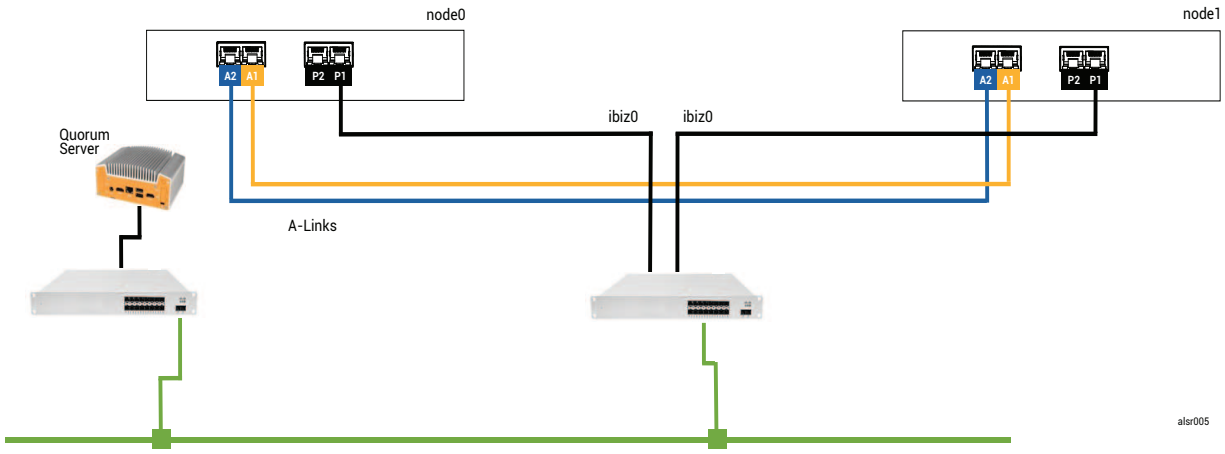

A Catastrophic Fault

That careless fork-truck operator crashes through the wall again, severing all of the network connections while leaving the power available and the system running. The following illustration shows the fault condition.

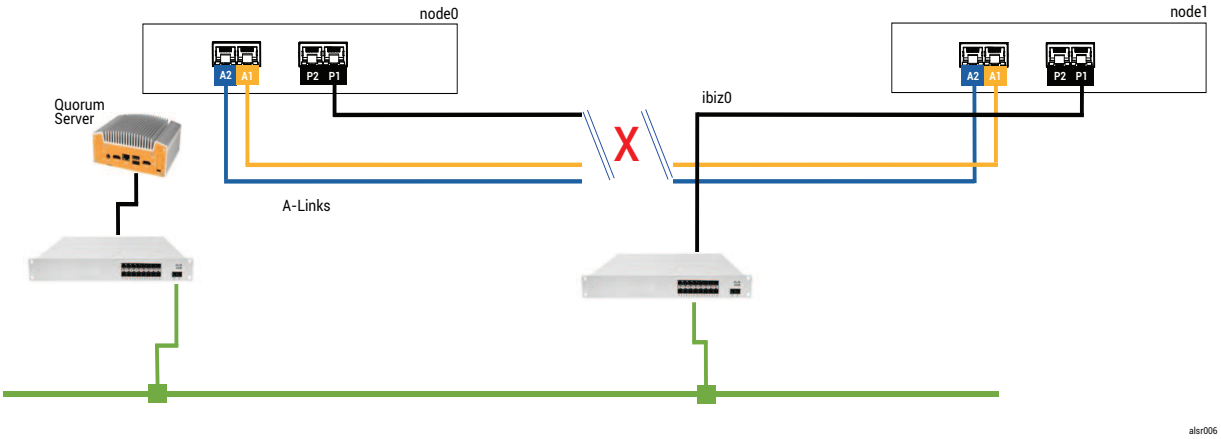

Fault Handling

The two nodes handle the fault, as follows:

- Node0—The AX on node0 detects the loss of both A-Links as well as all other network paths. Since the node0 AX can no longer detect the presence of its partner, the node0 AX attempts to contact the quorum server. In this case, the quorum server is also unavailable. Therefore, the node0 AX decides to shut down. The shutdown is not a graceful Windows shutdown, but is, instead, an abrupt stop, which causes the application inside the guest VM to stop.
- Node1—The AX on node1 also detects the loss of both A-Links, but ibiz0 remains available. The node1 AX tries to contact the quorum server, which responds, so the node1 AX remains active. The application inside the guest VM runs, perhaps not noticing any problems with the system.

> **Note**: If the node1 AX was not previously active and the guest VM is an HA VM, the guest VM on node1 might need to boot from node1's hard drive. In this case, the application experiences a brief period of downtime while the guest VM boots.

From the perspective of an application client or an external observer, the guest VM on node1 remains active and generates data while the VM on node0 is shut down. No split-brain condition exists.

Recovery and Repair

After some time, network connectivity is restored: the wall is repaired and the network cables are replaced.

When the node1 AX realizes that its partner is back online, the node0 AX becomes Standby. Because node0 was not previously running, data synchronization begins from node1 to node0.

Since a split-brain condition did not occur, no data is lost.

The system requires a few minutes to resynchronize, depending on how much disk activity needs to be sent to the standby node.

Example 2, Modified: The Quorum Server Is Unreachable During the Catastrophic Fault

In an ALSR system with a quorum server, the quorum server may be offline or otherwise unreachable when the catastrophic fault severs all of the network connections, though the power remains available and the system is still running. The following illustration shows a system in this situation with a quorum server that is offline.

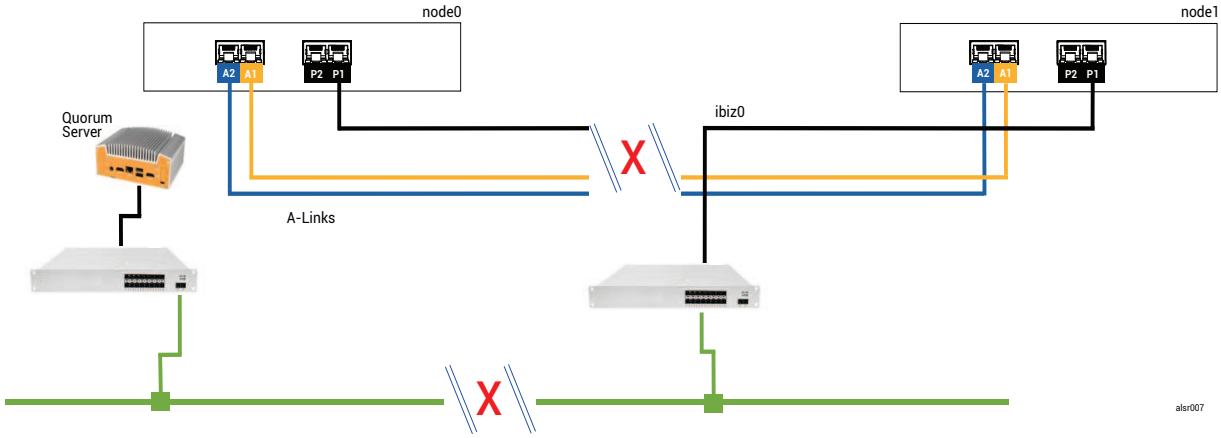

The fault handling is similar to Example 2 fault handling, with one important difference for node1:

> The node1 AX also detects the loss of both A-Links, but ibiz0 remains available. The node1 AX tries to contact the quorum server, but the communication fails. The AX terminates the guest VM.

In this case, the guest VM is shut down on both node0 and node1, preventing split-brain from occurring. The tradeoff is that the guest VM is unavailable until the connection to either node0 or to the quorum server is restored.

In this case, determine the node that you do not wish to operate and power it down. Then, forcibly boot the node that you wish to operate, and then forcibly boot the VM. For information on shutting down a VM and then starting it, see Managing the Operation of a Virtual Machine.)

Example 2, Modified: The Quorum Server Is Unreachable With No Catastrophic Fault

In some situations, the quorum server might be unreachable even without a catastrophic physical failure. One example is when the quorum computer is rebooted for routine maintenance such as applying an OS patch. In these situations, the AX detects that the quorum service is not responding and so the AX suspends synchronization traffic until the connection to the quorum server is restored. The guest VM continues to run on the node that was active when the connection was lost. However, the guest VM does not move to the standby node because additional faults may occur. After the quorum service is restored, the AX resumes synchronization and normal fault handling, as long as the connection to the quorum server is maintained.

Recovering From a Power Failure

If you restart the system after a power loss or a system shutdown, the ztC Edge system waits indefinitely for its partner to boot and respond before the system starts any guest VMs. If the AX that was previously active can contact the quorum server, the AX starts the guest VM immediately without waiting for the partner node to boot. If the AX that was previously standby boots first, it waits for its partner node.

If the system receives a response from either the partner node or the quorum server, normal operation resumes and the VM will start, subject to the same fault handler rules that apply in other cases.

If the system does not receive a response from the quorum server, or if the system does not have a quorum server, then a person must forcibly boot a guest VM, which overrides any decisions made by the AX or the fault handler. You must ensure that two people do not forcibly boot the same guest VM on node0 and node1. Doing so inadvertently causes a split-brain condition.

Accessing Knowledge Base Articles

The **Stratus Customer Service Portal** provides a searchable **Knowledge Base** with technical articles about all Stratus products, including ztC Edge. In some cases, the online Help directly references these Knowledge Base articles (for example, KB-*nnnn*). You can access the Customer Service Portal and Knowledge Base by using your existing service portal credentials, or by creating a new user account, as follows.

To access the Knowledge Base

1. Log on to the **Stratus Customer Service Portal** at https://support.stratus.com.

   If needed, create a new account as follows:

   a. Click **Register Account**.

   b. Enter your company email address and contact info, and click **Register**.

      Your company email address must include a domain name (for example, stratus.com) for a company that is a registered customer of Stratus.

   c. 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 Stratus service representative.

2. In the service portal, click **Knowledge Base** in the left pane.

3. In the **Keyword Search** box, enter keywords associated with the information you need, and then click **Search**.

   To search for an article by its KB-*nnnn* number, click **Advanced Search**. Next to **Search by ID**, type the article ID number (*nnnn*) and click **Display**.

Related Topics

Supporting Documents

REST API

Description

Get system information, including physical machine properties, statistics, system performance, and current alert list. The response can be large (about 14KB).

Header

Header	Value	Required
Locale	de (German), en-US (English), ja (Japanese), or zh-CN (Chinese). Default locale is en-US.	No
Content-type	application/json	Yes

Endpoint

```
GET /system/overview
```

Example

Request URL:

```
https://{hostname or IP address}/restapi/system/overview
```

10

Chapter 10: 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 an ztC Edge system to use SNMP, see Configuring SNMP Settings.

For information on using the `snmptable` command to obtain information about the system, specifically information about alerts, audit logs, nodes, VMs, and volumes, see Obtaining System Information with snmptable.

You can download a copy of the MIB file from the **Drivers and Tools** section of the **Downloads** page at https://www.stratus.com/services-support/downloads/?tab=ztcedge.

Obtaining System Information with snmptable

You can issue the `snmptable` command to obtain information about the system, specifically information about alerts, audit logs, nodes, VMs, and volumes.

To display alert information

To display information about alerts, issue the following command:

```
snmptable -v2c -m+/usr/smd/STRATUS-ZTC-EDGE-MIB.txt -c
public localhost ztCEdgeAlertTable
```

The command output displays the following:

Field	Description
`ztCEdgeAlertIndex`	The alert number.
`ztCEdgeAlertSeverity`	The alert severity (see `ztCEdgeAlertSeverityNum` for numerical value). Values are: `clear`  `informational`  `minor`  `major`  `serious`  `critical` 
`ztCEdgeAlertType`	The type of alert. Examples are: • `node_singleSystemDisk` • `Node Maintenance` • `The Unit is not well balanced`
`ztCEdgeAlertSource`	The source of the alert. Examples are: • `node0` or `node1` • ztC Edge system network name • network host name
`ztCEdgeAlertDateTime`	The date and time of the alert, in the format *yyyy-mm-dd hh:mm:ss*, where *yyyy* is year, *mm* is month, *dd* is date, *hh* is hour, *mm* is minute, and *ss* is second (for example, `2017-11-03 23:49:45`).
`ztCEdge AlertCallHomeSent`	If `true`, Call Home was sent; if `false`, it was not sent

Field	Description
ztCEdgeAlertEAlertSent	If true, e-Alert was sent; if false, it was not sent
ztCEdge AlertSNMPTrapSent	If true, SNMP trap was sent; if false, it was not sent
ztCEdgeAlertInformation	Information about the alert. Examples are: • Node node1 is in maintenance • node0 has a single system disk: Policy assumes this disk is redundant - if not, please add another internal disk • BUSINESS network net_728 is reporting a degraded link condition • The unit is not well load balanced
ztCEdgeAlertSNMPTrapOID	SNMP trap object identifier (OID) (for example, COMPANY-MIB::nodeSingleSystemDisk)
ztCEdgeAlertSeverityNum	ztCEdgeAlertSeverity number. Values are: 0 Clear 1 Informational 2 Minor 3 Major 4 Serious 5 Critical

To display audit log information

To display information about audit logs, issue the following command:

```
snmptable -v2c -m+/usr/smd/STRATUS-ZTC-EDGE-MIB.txt -c
public localhost ztCEdgeAuditTable
```

The command output displays the following:

Field	Description
`ztCEdgeAuditIndex`	An incrementing number (`1`, `2`, etc.) to indicate the audit log whose information is displayed..
`ztCEdgeAuditDateTime`	The date and time that the audit was generated, in the format *yyyy-mm-dd hh:mm:ss*, where *yyyy* is year, *mm* is month, *dd* is date, *hh* is hour, *mm* is minute, and *ss* is second (for example, `2017-11-03 23:49:45`).
`ztCEdgeAuditUsername`	The name of the user that caused the audit to be generated (for example, `audit` or `admin`).
`ztCEdge AuditOriginatingHost`	The IP address of the host that originated the audit.
`ztCEdgeAuditAction`	A description of the action being audited. Examples are: • `"Login user \"audit"` • `"Start virtual machine \"manager1"` • `"Remove all cleared alert"`

To display node information

To display node information, issue the following command:

```
snmptable -v2c -m+/usr/smd/STRATUS-ZTC-EDGE-MIB.txt -c
public localhost ztCEdgeNodeTable
```

The command output displays the following:

Field	Description
`ztCEdgeNodeIndex`	A number (typically `1` or `2`) to indicate the node whose information is displayed.
`ztCEdgeNodeId`	The host ID of the node (for example, `host:o34`).
`ztCEdgeNodeDisplayName`	The node name, `node0` or `node1`.
`ztCEdgeNodeIsPrimary`	If `true`, the node is primary. If `false`, the node is secondary.
`ztCEdgeNodeStateNum`	Node state is: `0` Normal () `1` Warning () `2` Busy () `3` Broken () `4` Maintenance ()
`ztCEdgeNodeActivityNum`	Node activity is: `0` Imaging `1` Booting `2` Running `3` Stopping `4` Rebooting `5` Powered off `6` Failed `7` Firmware updating `8` Lost `9` Exiled `10` Unreachable

Field	Description
	11 Proto (initializing) 12 Evacuating

To display VM information

To display VM information, issue the following command:

```
snmptable -v2c -m+/usr/smd/STRATUS-ZTC-EDGE-MIB.txt -c
public localhost ztCEdgeVMTable
```

The command output displays the following:

Field	Description
ztCEdgeVMIndex	An incrementing number (1, 2, etc.) to indicate the VM whose information is displayed.
ztCEdgeVMId	The VM ID (for example, vm:o1467).
ztCEdgeVMDisplayName	The VM name (for example, MyVM).
ztCEdgeVMRunningNode	The node on which the VM is running, node0 or node1.
ztCEdgeVMAvailability	The VM availability, HA (High Availability).
ztCEdgeVMStateNum	VM state is: 0 Normal () 1 Warning () 2 Busy or synchronizing () 3 Broken or blacklisted ()
ztCEdgeVMActivityNum	VM activity is: 0 Installing 1 Booting

Field	Description
	2 Running 3 Moving 4 Stopping 5 Stopped 6 Exporting 8 Paused 9 Loading 10 Crashing 11 Crashed 12 Dumping 13 Waiting

To display volume information

To display volume information, issue the following command:

```
snmptable -v2c -m+/usr/smd/STRATUS-ZTC-EDGE-MIB.txt -c
public localhost ztCEdgeVolumeTable
```

The command output displays the following:

Field	Description
ztCEdgeVolumeIndex	An incrementing number (1, 2, etc.) to indicate the volume whose information is displayed.
ztCEdgeVolumeId	The volume ID (for example, volume:o588).
ztCEdgeVolumeDisplayName	The volume name (for example, root).
ztCEdge VolumeSyncPercentage	The percentage of the volume that is synchronized.
ztCEdgeVolumeUsedBy	The name of the VM or host that is using the volume

Field	Description
	(for example, `MyVM`); `none` indicates that the volume is not being used.
`ztCEdgeVolumeStateNum`	Volume state is: `0` Normal () `1` Warning () `2` Busy or synchronizing () `3` Broken ()