

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 Bermuda, 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.

This document is protected by copyright. All rights are reserved. Stratus Technologies grants you limited permission to download and print a reasonable number of copies of this document (or any portions thereof), without charge, for your internal use only, provided you retain all copyright notices and other restrictive legends and/or notices appearing in the copied document.

Copyrights

Stratus, the Stratus logo, everRun, and SplitSite are registered trademarks of Stratus Technologies Bermuda, Ltd. The Stratus Technologies logo, the Stratus 24 x 7 logo, and Automated Uptime are trademarks of Stratus Technologies Bermuda, Ltd.

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

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

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

VMware is a registered trademarks of VMware, Inc. in the United States and/or other jurisdictions.

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

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

Mozilla and Firefox are registered trademarks of the Mozilla Foundation.

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

Dell is a trademark of Dell Inc.

Hewlett-Packard and HP are registered trademarks of Hewlett-Packard Company.

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

Manual Name: everRun Enterprise User's Guide

Product Release Number: everRun Enterprise Release 7.1.0.0

Publication Date: Tuesday, July 22, 2014

Stratus Technologies, Inc.

111 Powdermill Road

Maynard, Massachusetts 01754-3409

© 2014 Stratus Technologies Bermuda, Ltd. All rights reserved.

Table of Contents

Part 1: everRun Enterprise User's Guide	1
Chapter 1: Introduction to everRun Enterprise Systems	1
everRun Enterprise Quick Start Guide	1
Assembling Required Items	2
Configuring Your RAID Controller	2
Cabling Your System	3
Burning the Software to a DVD	4
Installing the everRun Enterprise Software	5
Logging On to the everRun Availability Console	6
Creating a Protected Virtual Machine	7
everRun Enterprise System Overview	7
everRun Enterprise System Description	8
Physical Machines and Virtual Machines	8
Administrative Operations	9
Alerts	9
Remote Support	10
Lights-Out Management	10
Third-party Management Tools	10
Modes of Operation	11
High Availability Operation	12
Fault Tolerant Operation	12
SplitSite Configurations	13
SplitSite and Quorum Service	13
Quorum Servers	14
everRun Enterprise Storage Architecture	15
Logical Disks and Physical Disks	15
The Storage Group	15
Network Architecture	16
Network Architecture Overview	16
A-Link and Private Networks	16
Business and Management Networks	17
System Usage Restrictions	18

QEMU	18
Accessing the Host Operating System	18
Chapter 2: Getting Started	21
Planning	21
System Requirements Overview	22
System Hardware	22
Supported Servers	22
RAM	22
Disk Space Requirements	22
Network	22
IP Addresses	23
Ports	23
System Software	23
Storage Requirements	23
Memory Requirements	24
General Network Requirements and Configurations	24
Requirements	24
Recommended Configurations	25
Business and Management Network Requirements	25
A-Link and Private Network Requirements	26
SplitSite Network Requirements	27
A-Link Network Requirements	27
Private Network Requirements	28
Business Network Requirements	28
Management Network Requirements	29
everRun Enterprise Availability Console Requirements	29
Compatible Internet Browsers	29
Java™ Requirements	30
Quorum Servers Considerations	30
Power Requirements and Considerations	31
Software Installation	32
Site and System Preparation	32
Connecting Power	33
UPS (Optional)	33

Obtaining everRun Enterprise Software	34
Configuring the BIOS	35
Required Settings	35
Recommended Settings	36
Installing everRun Enterprise Software	36
Connecting Ethernet Cables	37
Installation Options	38
Installing Software on the First PM	39
Recording the Management IP Address	43
Installing Software on the Second PM	43
Post-Installation Tasks	45
Obtaining System IP Information	46
Logging on to the everRun Availability Console for the First Time	46
Connecting Additional Networks	48
Chapter 3: Using the everRun Availability Console	51
The everRun Availability Console	51
Logging on to the everRun Availability Console	53
The Dashboard Page	53
Resolving Outstanding Alerts on the Dashboard	54
The System Page	55
Rebooting the System	55
Shutting Down the System	56
The Preferences Page	57
Specifying Owner Information	59
Managing the everRun Enterprise Product License	60
Configuring IP Settings	64
Configuring Quorum Servers	67
Configuring Date and Time	68
Configuring System Resources	69
Configuring the Virtual Machine Import Option	70
Managing Diagnostic Files	71
Creating a Diagnostic File	71
Uploading a Diagnostic File to Customer Support	72
Deleting a Diagnostic File	73

Configuring e-Alerts	73
Configuring SNMP Settings	75
Configuring Remote Support Settings	77
Configuring Internet Proxy Settings	78
The Alerts Page	79
The Audits Page	80
The Physical Machines Page	80
Physical Machine Actions	81
Physical Machine States and Activities	83
The Virtual Machines Page	84
Virtual Machine Actions	85
Virtual Machine States and Activities	87
The Volumes Page	88
The Storage Group Page	89
The Networks Page	90
The Virtual CDs Page	90
The Upgrade Kits Page	91
The Users & Groups Page	91
Managing Local User Accounts	92
User Roles	93
Chapter 4: Upgrading everRun Enterprise Software	95
Chapter 5: Migrating From Non-everRun Enterprise Systems	97
Planning to Migrate from an everRun MX System	97
Platform Requirements	98
Planned Outage	98
Guest Operating System Support	98
Network Preparation	98
Management Network Access	98
Availability Link Networks	99
Private Network	99
Business Networks	99
Storage Considerations	99
Quorum Support	100
Installing everRun Enterprise	100

Migrating Virtual Machines	100
Planning to Migrate from an Avance Unit	100
Platform Requirements	100
Planned Outage	101
Guest Operating System Support	101
Network Preparation	101
Management Network Access	101
Availability Link Networks	101
Private Network	102
Business Networks	102
Storage Considerations	102
Installing everRun Enterprise	102
Migrating Virtual Machines	102
Chapter 6: Managing Logical Disks	103
Logical Disk Management	103
Responding to a Failed Logical Disk	104
Activating a New Logical Disk	106
Chapter 7: Managing Physical Machines	109
Maintenance Mode	109
Physical Machine Management Actions	111
Rebooting a Physical Machine	111
Shutting Down a Physical Machine	112
Load Balancing	113
Modes of Operation	113
Troubleshooting Physical Machines	113
Recovering a Failed Physical Machine	114
Resetting MTBF for a Failed Physical Machine	116
Chapter 8: Managing Virtual Machines	119
Planning Virtual Machine Resources	120
Planning Virtual Machine vCPUs	120
Planning Virtual Machine Memory	122
Planning Virtual Machine Storage	123
Planning Virtual Machine Networks	124
Creating and Migrating Virtual Machines	125

Creating a New Virtual Machine	125
Migrating a Physical Machine or Virtual Machine to the everRun Enterprise System	128
Importing an OVF File from an everRun MX System to the everRun Enterprise System	135
Importing an OVF File from an Avance System to the everRun Enterprise System	143
Managing Windows Drive Labels	151
Configuring Windows-based Virtual Machines	152
Creating and Initializing a Disk (Windows-based VMs)	153
Installing Applications (Windows-based VMs)	153
Configuring Linux-based Virtual Machines	154
Creating and Initializing a Disk (Linux-based VMs)	155
Installing Applications (Linux-based VMs)	155
Managing the Operation of a Virtual Machine	156
Starting a Virtual Machine	156
Shutting Down a Virtual Machine	156
Powering Off a Virtual Machine	157
Opening a Virtual Machine Console Session	158
Renaming a Virtual Machine	159
Removing a Virtual Machine	160
Managing Virtual Machine Resources	160
Reprovisioning Virtual Machine Resources	161
Creating a Volume in a Virtual Machine	163
Attaching a Volume to a Virtual Machine	164
Detaching a Volume from a Virtual Machine	165
Removing a Volume from a Virtual Machine	166
Renaming a Volume on the everRun System	167
Recovering Virtual Machine Resources	168
Managing Virtual CDs	169
Creating a Virtual CD	169
Burning a CD or DVD for a Virtual CD	171
Booting from a Virtual CD	171
Renaming a Virtual CD	172
Removing a Virtual CD	172
Advanced Topics (Virtual Machines)	173
Assigning a Specific MAC Address to a Virtual Machine	173

Selecting a Preferred PM for a Virtual Machine	174
Changing the Protection Level for a Virtual Machine (HA or FT)	174
Configuring the Boot Sequence for Virtual Machines	175
Resetting MTBF for a Failed Virtual Machine	176
Locating a Dump File in a Virtual Machine	177
Chapter 9: Maintaining Physical Machines	179
Physical Machine Hardware Maintenance Restrictions	179
Adding or Replacing Hot-Swappable Components	180
Adding or Replacing Components That Are Not Hot-Swappable	181
Replacing Physical Machines, Motherboards, NICs, or RAID Controllers	182
Upgrading Both Physical Machines In a Running System	185
Part 2: Supporting Documents	187
Chapter 10: everRun Enterprise Release 7.1.0.0 Release Notes	188
Important Considerations	188
Upgrading From Previous Releases of everRun Enterprise	188
Guest Volumes Must be Smaller Than 2 TB	189
The Active Directory Feature is Not Supported	189
Optimizing Performance of A-Link Networks	189
Migrating a PM or VM to an everRun Enterprise System	190
Status of RAID Physical Disks is Not Monitored	190
Other Important everRun Enterprise Considerations	190
Known Issues	190
Replacing Physical Machines, Motherboards, NICs, or RAID Controllers	190
Error Message May Appear When Importing a VM to a HA-only Licensed System	191
Using Internet Explorer Can't Change Volume Name During VM Import or Restore	192
Unsupported Network Adapter Card and Chip	192
Do Not Use the ifdown Command	192
New Features, Enhancements, and Bug Fixes	192
New in everRun Enterprise Release 7.1.0.0	192
Getting Help	193
Chapter 11: everRun Enterprise Command Line Interface Reference	194
AVCLI Command Overview	194
Prerequisites	195
Installing the Linux Client	195

Installing the Windows Client	196
Using AVCLI	197
Executing a Command	197
Using AVCLI Help	198
Listing All Commands	198
Displaying Help for a Specific Command	199
AVCLI Error Status	199
XML Encapsulated Errors	199
Error Checking	200
Asynchronous Command Delay	200
Output Formatting	201
User-Friendly Command Output	201
Program-Friendly XML Output	202
AVCLI Exceptions	205
AVCLI Command Descriptions	206
alert-delete	211
alert-info	212
audit-export	213
audit-info	214
callhome-disable	215
callhome-enable	216
callhome-info	217
datetime-config	218
diagnostic-create	221
diagnostic-delete	222
diagnostic-extract	223
diagnostic-fetch	224
diagnostic-info	225
dialin-disable	226
dialin-enable	227
dialin-info	228
ealert-config	229
ealert-disable	230
ealert-enable	231

ealert-info	232
help	233
kit-delete	234
kit-info	235
kit-upload	236
license-info	237
license-install	238
local-group-add	239
local-group-delete	240
local-group-edit	241
local-group-info	242
local-user-add	243
local-user-delete	245
local-user-edit	246
local-user-info	248
media-create	249
media-delete	250
media-eject	251
media-import	252
media-info	254
network-change-mtu	255
network-change-role	256
network-info	257
node-add	259
node-cancel	260
node-delete	261
node-info	262
node-poweroff	263
node-poweron	264
node-reboot	265
node-recover	266
node-shutdown	267
node-upgrade	268
node-workoff	269

node-workon	270
ntp-config	271
ntp-disable	272
ova-info	273
ovf-info	274
owner-config	275
owner-info	276
proxy-config	277
proxy-disable	278
proxy-enable	279
proxy-info	280
snmp-config	281
snmp-disable	282
snmp-info	283
storage-group-info	284
storage-info	285
timezone-config	286
timezone-info	287
unit-configure	288
unit-eula-accept	289
unit-eula-reset	290
unit-info	291
unit-shutdown	292
unit-shutdown-cancel	293
unit-shutdown-state	294
unit-synced	295
vm-boot-attributes	296
vm-cd-boot	297
vm-create	298
vm-delete	301
vm-import	302
vm-info	305
vm-migrate	306
vm-poweroff	307

vm-poweron	308
vm-reprovision	309
vm-shutdown	312
vm-unlock	313
volume-info	314
Chapter 12: System Reference Information	316
Compatible Guest Operating Systems	316
Physical Machine System Requirements	317
Important Physical Machine and Virtual Machine Considerations	319
Virtual Machine Recommendations and Limits	319
Recommended Number of CPU Cores	319
Virtual Machine Limits	320
Combined Virtual Machine Maximums	321
Important Considerations	321
Chapter 13: SNMP	324
MIB File Contents	324

Part 1: everRun Enterprise User's Guide

The everRun Enterprise User's Guide describes everRun Enterprise systems, how to install them, and how to use them.

For a summary of the steps required to install everRun Enterprise software, see:

- ["everRun Enterprise Quick Start Guide" on page 1](#)

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

- ["Introduction to everRun Enterprise Systems" on page 1](#)

For planning and installation information, see:

- ["Getting Started" on page 21](#)

The following topics describe how to administer everRun Enterprise systems.

- ["Using the everRun Availability Console" on page 51](#)
- ["Upgrading everRun Enterprise Software" on page 95](#)
- ["Migrating From Non-everRun Enterprise Systems" on page 97](#)
- ["Managing Logical Disks" on page 103](#)
- ["Managing Physical Machines" on page 109](#)

- ["Managing Virtual Machines" on page 119](#)
- ["Maintaining Physical Machines" on page 179](#)

1

Chapter 1: Introduction to everRun Enterprise Systems

For a summary of the steps required to install everRun Enterprise software, see the ["everRun Enterprise Quick Start Guide" on page 1](#).

See the following topics for an introduction to everRun Enterprise systems:

- ["everRun Enterprise System Overview" on page 7](#)
- ["Modes of Operation" on page 11](#)
- ["everRun Enterprise Storage Architecture" on page 15](#)
- ["Network Architecture" on page 16](#)
- ["System Usage Restrictions" on page 18](#)

everRun Enterprise Quick Start Guide

Use this Quick Start Guide to get your everRun Enterprise system up and running as quickly as possible.

An everRun Enterprise system requires two x86-64 host servers (referred to as *physical machines*, or *PMs*) that can support multiple virtual machines (VMs), and a remote management computer that can run the everRun Availability Console. This guide explains how to set up your PMs and guides you through the basic installation and start-up tasks, including:

- ["Assembling Required Items" on page 2](#)
- ["Configuring Your RAID Controller" on page 2](#)
- ["Cabling Your System" on page 3](#)
- ["Burning the Software to a DVD" on page 4](#)

- ["Installing the everRun Enterprise Software" on page 5](#)
- ["Logging On to the everRun Availability Console " on page 6](#)
- ["Creating a Protected Virtual Machine" on page 7](#)

Note: If you need assistance during the installation process:



- Call 866-763-1813 (U.S. Toll-Free) or 602-852-3094 (International)
- Visit the **everRun Enterprise Downloads and Support** page at <http://www.stratus.com/go/support/everrun>

Assembling Required Items

You need the following items/information:

- Two PMs that meet the requirements listed in ["System Requirements Overview" on page 22](#)
- Ethernet cables for each network you are connecting
- A remote management computer. This is a general-purpose PC with a supported web browser to access the everRun Availability Console. It must be located on the same business/management network as the PMs being installed. For details, see ["everRun Enterprise Availability Console Requirements" on page 29](#).
- A monitor, keyboard, and cables, used only during the installation
- The everRun Enterprise license key you received from Stratus
- The everRun Enterprise ISO image available for download at the **everRun Enterprise Downloads and Support** page at <http://www.stratus.com/go/support/everrun>
- From your network administrator, the IPv4 address, netmask, default gateway address, and DNS address values for everRun Enterprise and each PM

Configuring Your RAID Controller

Stratus strongly recommends that your everRun Enterprise system uses a storage RAID controller. The RAID controllers in an everRun Enterprise system create logical disks from the system's physical disks, and the logical disks are then collected into a storage group. Configuration recommendations follow:

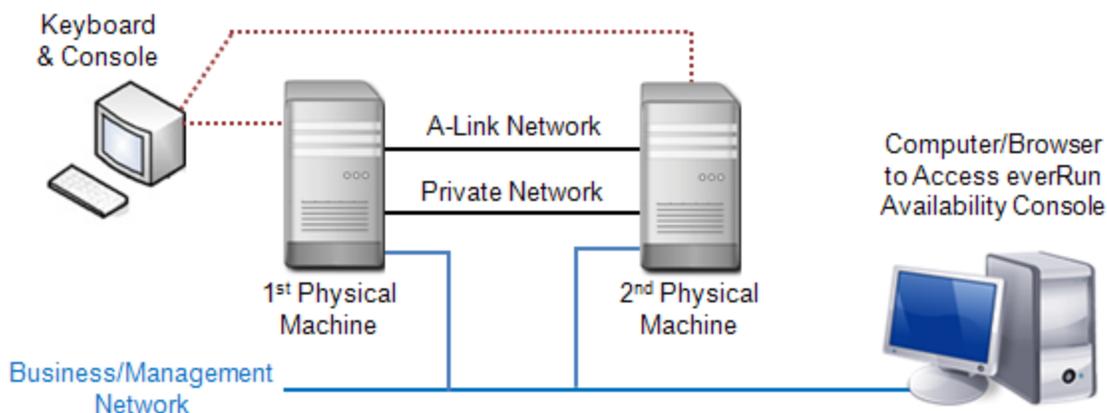
- If the system has a single logical disk, Stratus strongly recommends that you configure the RAID controller so that logical disks presented to the host are backed by redundant physical drives.
- Stratus strongly recommends that RAID controllers have a battery-backed write cache.
- You must configure the RAID controller to boot off the first logical disk.

Cabling Your System

Attach the following cables:

- Private network: Connect an Ethernet cable from the **first embedded port** on the first PM to the **first embedded port** on the second PM. If you plan to use the private network as an A-Link, see ["A-Link and Private Networks" on page 16](#).
- Business/Management network: The first business network is the *management network*. Connect Ethernet cables from the **second embedded port** on each PM, through a network switch to a network, and connect the remote management computer to this network.
- A-Link network(s): For each A-Link network, connect an Ethernet cable from any unused port on the first PM to any unused port on the second PM, either directly or through a network switch.
- Business network(s): For each business network, connect Ethernet cables from a port on the first PM to a port on the second PM, through a network switch to a network.
- Make sure the remote management computer is connected, or routes to, the management network.
- Connect the monitor, keyboard, and mouse to the first PM. See ["Site and System Preparation" on page 32](#) for details.

The following illustration shows these connections:





Note: When installing software on the first PM, connect the keyboard and monitor to the first PM. When installing software on the second PM, connect the keyboard and monitor to the second PM. When software installation is complete, disconnect the keyboard and monitor from the system.

Burning the Software to a DVD

Obtain the ISO image, validate it, and burn it to a DVD:

1. From any computer connected to the Internet, go to the **everRun Enterprise Downloads and Support** page at <http://www.stratus.com/go/support/everrun>.
2. To download the everRun Enterprise software ISO image (**everRun_Enterprise_install-7.x.x.x-xx.iso**), under **Product Download**, click **everRun Enterprise 7.x.x.x ISO Image**. Save the ISO image.

Occasionally a file can be corrupted during the download process. To verify that the downloaded file is not corrupted, perform steps 3 through 7, otherwise skip to step 8.

3. To download the **MSUMS** file, under **Product Download**, click **everRun Enterprise 7.x.x.x md5sum**. Save the file. This file is used to validate the downloaded ISO image to make sure that it was not corrupted during the download.
4. From the Internet, find the **md5sum.exe** file and download it to the directory containing the everRun Enterprise ISO image.
5. If you have a Windows-based system, open a command prompt.
6. Go to the directory containing the **md5sum.exe**, **md5**, and ISO image files, and then type the following command to check the ISO image's status:
 - On Windows: **md5sum.exe -c everRun_Enterprise_install-7.x.x.x.xxx.md5**
 - On Linux: **md5sum -c everRun_Enterprise_install-7.x.x.x.xxx.md5**
7. If the command succeeds (that is, it returns the message `everRun_Enterprise_install-7.x.x.x.-xxx.iso: OK`), go to the next step. If the command fails, repeat the download.

8. When the validation is complete, use a commonly available DVD application to burn the ISO image to a DVD. For example, if you have the Roxio application installed, right click the ISO file and select the option to burn a DVD.

For additional information, see ["Obtaining everRun Enterprise Software" on page 34](#).

Installing the everRun Enterprise Software

Allow 60 to 90 minutes to complete the everRun Enterprise software installation process.

1. Install the everRun Enterprise software on the first PM:
 - a. Power on the first PM and insert the DVD.
 - b. As the PM powers on, configure the following BIOS settings:
 - Set the first boot device to the Optical Drive.
 - Enable Virtualization Technology.
 - Enable Execute-Disable Bit Capability.
 - c. From the installation software's Welcome screen, use the arrow keys to select **Install everRun, Create a new system**, and press **Enter**.
 - d. From the **Select interface for private Physical Machine connection** screen, select the first embedded port, **em1** (if it is not already selected), and press **F12**.
 - e. From the **Select interface for managing the system (ibiz0)** screen, select the second embedded port, **em2** (if it is not already selected), and press **F12**.
 - f. From the **Select the method to configure ibiz0** screen, select **Manual configuration (Static Address)** and press **F12**.



Note: To perform a dynamic IP configuration, select **Automatic configuration via DHCP** instead and skip to **Step 1h**, where you will need to write down the IPv4 address as described in ["Recording the Management IP Address" on page 43](#).

- g. From the **Configure em2** screen, enter the IPv4 address, netmask, default gateway address, and DNS address values that you received from your network administrator, and then press **F12**.

6. The **Portal Restart Required** window appears. After waiting a minute (as indicated in the window), click **OK** to refresh the console and to continue.
7. The **LICENSE INFORMATION** window appears. Under **Upload License Key**, click **Browse** and navigate to the license **.KEY** file that you received from Stratus. Select the license file and click **Upload**.

For security, change the default user login name and password for the **admin** account on the **Users & Groups** page.

The everRun Availability Console appears. Bookmark or make note of the system IP address for use when logging in to the console in the future.

Creating a Protected Virtual Machine

First, create a virtual CD (VCD) to make software installation media available to the virtual machine (VM).

1. Open the **Virtual CDs** page in the everRun Availability Console
2. Click **Create VCD** to open the **Virtual CD Creation Wizard**.
3. Follow creation wizard prompts. For details, see ["Creating a Virtual CD" on page 169](#) in the online help.

Next, create a new virtual machine (VM) and install a guest operating system on your everRun Enterprise system.

1. On the **Virtual Machines** page, click **Create** to open the **VM Creation Wizard**.
2. Follow the creation wizard prompts. For details, see ["Creating a New Virtual Machine" on page 125](#) in the online help.

After you install the operating system, perform any other guest OS configuration tasks (for example, initializing disks and installing applications). For details, see ["Post-Installation Tasks" on page 45](#) in the online help.

everRun Enterprise System Overview

An everRun Enterprise system provides uninterrupted operation with no lost data in the event of a hardware failure. See the following topics for descriptions of system features and capabilities.

- ["everRun Enterprise System Description" on page 8](#)
- ["Physical Machines and Virtual Machines" on page 8](#)

- ["Administrative Operations" on page 9](#)
- ["Alerts" on page 9](#)
- ["Remote Support" on page 10](#)
- ["Lights-Out Management" on page 10](#)
- ["Third-party Management Tools" on page 10](#)

everRun Enterprise System Description

everRun Enterprise software allows two computers to work as a single, highly-available or fault-tolerant system. Each computer is called a physical machine.

Both physical machines (PMs):

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

The PMs must:

- Have compatible CPUs
- Conform to hardware requirements for everRun Enterprise systems. See ["Physical Machine System Requirements" on page 317](#) and ["System Requirements Overview" on page 22](#) for more information.

The data and memory contents of the two PMs are synchronized via direct Ethernet links. Other Ethernet connections to a network support virtual machine and management operations.

Related Topics

["System Requirements Overview" on page 22](#)

["Compatible Guest Operating Systems" on page 316](#)

["Network Architecture Overview" on page 16](#)

Physical Machines and Virtual Machines

An everRun Enterprise system transparently protects applications by creating redundant virtual machines (VMs) that run on two physical machines (PMs).

The everRun management software can create an everRun-protected VM from scratch. It is also possible to import existing VMs from other environments and convert them into an everRun-protected VMs. By creating an identical instance of the selected VM on a second host PM, everRun software provides FT-class protection of the VM. The system administrator manages this single entity from a separate, browser-based management console called the everRun Availability 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 protected 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, everRun software automatically redirects I/O to the paired host PM for continuous operation. Though redundancy is lost until the failure is repaired, the client experiences no interruption in connectivity and no loss of data. 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, everRun software automatically incorporates the repaired components into the protected environment and restores redundancy without interrupting the application.

Related Topics

["Using the everRun Availability Console" on page 51](#)

["The Physical Machines Page" on page 80](#)

["The Virtual Machines Page" on page 84](#)

Administrative Operations

You can perform many administrative operations on the everRun Enterprise system from the everRun Availability 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 everRun Availability Console" on page 51](#).

Alerts

everRun Enterprise 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 everRun Enterprise 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](#)" on page 77 for details.
- **Proxy Configuration**—Enables you to configure a proxy server for access to the Internet. See "[Configuring Internet Proxy Settings](#)" on page 78 for details.

Lights-Out Management

Some server vendors may provide lights-out capabilities. Lights-out capabilities enable administrators to perform a wide range of system management and operation functions remotely. everRun Enterprise systems fully support lights-out management on vendor servers.

Third-party Management Tools

You can install third-party management tools on everRun Enterprise systems. Examples of such tools include vendor- or platform-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 everRun systems. Possible exceptions are tools that manage/monitor the CentOS KVM-based virtualization. To manage/monitor everRun virtualization, use the integrated everRun management tools.
- Before deploying your everRun system, Stratus recommends that you verify that it operates properly with the management tools installed and operational.
- You can access your everRun system by using the root password specified during the system installation process. Stratus recommends that you set up a non-root account for third-party management tools.
- You can access your everRun 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).

For information about accessing the host operating system, see ["Accessing the Host Operating System" on page 18](#).

Related Topics

["Getting Started" on page 21](#)

["System Reference Information" on page 316](#)

Modes of Operation

An everRun Enterprise system provides two modes of operation to set user-defined availability levels for VMs:

- ["High Availability Operation" on page 12](#)
- ["Fault Tolerant Operation" on page 12](#)

Both HA operation and FT operation achieve their respective level of redundancy by using a pair of physical machines (PMs).

Stratus recommends configuring quorum service for both HA operation and FT operation. The quorum service prevents a condition called *split-brain* where both PMs of an HA operation and FT operation pair are running independently of each other; for information, see ["Quorum Servers" on page 14](#).

High Availability Operation

everRun Enterprise software provides two user-defined availability levels for VMs: High Availability (HA) and Fault Tolerant (FT).

In HA operation, everRun Enterprise software automatically detects, isolates, and handles most hardware faults, thereby keeping your applications running. With HA remote-support technology, the everRun 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.

You select a VM's availability level when you create or import the VM by using the everRun Availability Console.

When enabled, 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" on page 84](#)

["Using the everRun Availability Console" on page 51](#)

Fault Tolerant Operation

everRun Enterprise software provides two user-defined availability levels for VMs: High Availability (HA) and Fault Tolerant (FT). In FT operation, an application continues to run without downtime during a fault. Use FT for applications that need the highest levels of availability.

You select a VM's availability level when you create or import the VM by using the everRun Availability Console.

In FT operation, the everRun software transparently protects an application by creating a redundant environment for a VM running across two physical machines (PMs). With an identical instance of the selected VM on a second host, the everRun software provides FT-class protection of the VM.

When enabled, FT operation transparently protects a VM from all faults, with no downtime, and FT operation:

- Eliminates downtime due to any CPU, memory, I/O, or other physical machine (PM) failure.
- Handles failures without IT intervention.
- Ensures no data loss.
- Provides continuous, active validation of all components.
- Assures complete redundancy and recovery at all times.

Related Topics

["The Virtual Machines Page" on page 84](#)

["Using the everRun Availability Console" on page 51](#)

SplitSite Configurations



Note: An everRun Enterprise SplitSite license is required to run SplitSite configurations.

A *SplitSite 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, a SplitSite configuration requires careful planning of component placement and more complex networking topologies. **For SplitSite configurations, Stratus strongly recommends that you use the quorum service because a SplitSite configuration exposes the A-Link networks to other potential failure scenarios.**

["SplitSite Network Requirements" on page 27](#) lists the requirements for networks in a SplitSite configuration.

SplitSite and Quorum Service

In a SplitSite configuration, configure two quorum-service computers in compliance with the best practices recommended for quorum deployment (see ["Quorum Servers Considerations" on page 30](#)). In any SplitSite 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 everRun Enterprise 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 everRun Enterprise 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 everRun Enterprise PMs and the quorum-service computers. For this to work properly, you must configure each everRun Enterprise 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

["Quorum Servers" on page 14](#)

["Network Architecture Overview" on page 16](#)

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) running HA- or FT-protected virtual machines (VMs). Quorum servers provide data integrity assurances and automatic restart capabilities for specific failures in an everRun Enterprise environment. Stratus strongly recommends using quorum servers, especially for SplitSite operation. You can configure an everRun Enterprise PM pair 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 SplitSite configurations. Best practice for SplitSite 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.

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 a potential split-brain scenario.

Related Topics

["Quorum Servers Considerations" on page 30](#)

["Configuring Quorum Servers" on page 67](#)

["SplitSite Configurations" on page 13](#)

everRun Enterprise Storage Architecture

The RAID controllers in an everRun Enterprise system create logical disks from the system's physical disks. The logical disks are collected into a storage group. For more information about everRun Enterprise storage architectures, see the following topics:

- ["Logical Disks and Physical Disks" on page 15](#)
- ["The Storage Group" on page 15](#)

Logical Disks and Physical Disks

In an everRun Enterprise system, the RAID controller creates logical disks from the system's physical disks. everRun software is able to access logical disks that the RAID controller presents to the operating system. The everRun software detects new logical disks and logical disk failures. You manage logical disks using the everRun Availability Console. For information, see ["Managing Logical Disks" on page 103](#).

You need to use the RAID controller to manage and monitor physical disks. Follow the RAID controller manufacturer's requirements to add a new or replacement physical disk to a RAID array.

Related Topics

["The everRun Availability Console" on page 51](#)

["Using the everRun Availability Console" on page 51](#)

The Storage Group

In an everRun Enterprise system, a storage group is a collection of logical disks. everRun Enterprise software creates the **Initial Storage Group**, which contains all logical disks present at install time. You can view information about the storage group from the **Storage Groups** page of the everRun Availability Console. For information, see ["The Storage Group Page" on page 89](#).

Related Topics

["The everRun Availability Console" on page 51](#)

["Using the everRun Availability Console" on page 51](#)

Network Architecture

See the following topics for information about everRun Enterprise network architecture.

- ["Network Architecture Overview" on page 16](#)
- ["A-Link and Private Networks" on page 16](#)
- ["Business and Management Networks" on page 17](#)

Network Architecture Overview

Ethernet networks provide pathways for communications between two physical machines (PMs) in an everRun Enterprise system. The main types of Ethernet networks are:

- *Availability Link networks*, or *A-Link networks*, are assigned to virtual machines (VMs) and are used to synchronize data or migrate VMs between two PMs. One A-Link network must be a *private network*, which connects the two everRun PMs. See ["A-Link and Private Networks" on page 16](#).
- *Business networks* allow your applications to connect to your network. One business network must be a *management network*, which connects to the everRun Availability Console and is used by the quorum servers. See ["Business and Management Networks" on page 17](#).

An everRun system must have a minimum of one private network and one management network per PM.

A-Link and Private Networks

Every everRun Enterprise system requires one *private* network, referred to as *priv0*, that connects the two everRun physical machines (PMs). The private network is used only for discovery and can have no other entities on it that respond to IPv4 broadcasts.

In addition to the private network, an everRun system has A-Link networks to increase data-replication performance between PMs. *A-Link* networks let you sync disks, shunt networks, migrate, perform heartbeat checks, and sync fault-tolerant memory.

By default, the private network also performs the role of an A-Link network under the following conditions:

- If the private network speed is at least 10 Gb
- If the private network speed is less than 10 Gb and the system contains no other 10 Gb ports (other than the management link). In this situation, you can remove the A-link role later as long as the private network is not currently in use as an A-Link **and** is not the only remaining A-Link.

The private network cannot perform the A-Link role if its speed is less than 10 Gb **and** the system contains any 10 Gb ports (other than the management link). However, you can assign the A-Link role to the private network at a later time.

The simplest private network 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 the private network, see ["SplitSite Configurations" on page 13](#).

Connect A-link networks between PMs either directly (that is, in the same manner that you connect the private network) or through a network switch.

Be sure to set up redundant A-Link networks.

The everRun Enterprise installation software sets up the private network. It also sets up A-Link networks for any A-Link network ports that are physically connected at the time of the installation. To set up an A-Link network after the installation is complete (recommended if the network includes many additional A-Link network ports), see ["Connecting Additional Networks" on page 48](#).

Related Topics

["Business and Management Networks" on page 17](#)

["A-Link and Private Network Requirements" on page 26](#)

["Network Architecture Overview" on page 16](#)

Business and Management Networks

All Ethernet ports--other than those used by A-Link networks (including the private network port)--are considered business-network ports, which your guest operating systems use to connect to your network.

One business network is the *management* network, which accesses the everRun Availability Console and handles miscellaneous management tasks and the quorum server. Each everRun Enterprise PM has a single management network, referred to as *ibiz0*.

The everRun Enterprise installation software sets up the management network. It also sets up business networks for any business-network ports that are physically connected at the time of the installation. To set up business networks after the installation is complete, see ["Connecting Additional Networks" on page 48](#).

Related Topics

["A-Link and Private Networks" on page 16](#)

["Business and Management Network Requirements" on page 25](#)

["Network Architecture Overview" on page 16](#)

System Usage Restrictions

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

- ["QEMU" on page 18](#)
- ["Accessing the Host Operating System" on page 18](#)

QEMU

Stratus everRun Enterprise 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.

everRun Enterprise users should make no changes to the QEMU virtualization engine or to its configuration.

Accessing the Host Operating System

After the everRun software installation process completes, you can access the host operating system (CentOS) locally via the PM's physical console, or you can access it remotely via SSH.

If you access the host operating system via SSH, use the management IP addresses specified during the installation process (or supplied by the DHCP server, if the interface was configured for DHCP during installation). See ["Recording the Management IP Address" on page 43](#).



Note: Do not use the system IP address to access the host operating system, as it can move from PM to PM.

The password for the root account is **everRun**.



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" on page 10](#).

2

Chapter 2: Getting Started

The following topics describe the everRun Enterprise planning, installation and post-installation tasks:

- ["Planning" on page 21](#)
- ["Software Installation" on page 32](#)
- ["Post-Installation Tasks" on page 45](#)

Planning

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

- ["System Requirements Overview" on page 22](#)
- ["Storage Requirements" on page 23](#)
- ["Memory Requirements" on page 24](#)
- ["General Network Requirements and Configurations" on page 24](#)
- ["Business and Management Network Requirements" on page 25](#)
- ["A-Link and Private Network Requirements" on page 26](#)
- ["SplitSite Network Requirements" on page 27](#)
- ["everRun Enterprise Availability Console Requirements" on page 29](#)
- ["Compatible Internet Browsers" on page 29](#)
- ["Quorum Servers Considerations" on page 30](#)
- ["Power Requirements and Considerations" on page 31](#)

System Requirements Overview

An everRun Enterprise system requires two x86-64 host servers that can support multiple virtual machines (VMs), and a remote management computer (that is, a general-purpose PC) that can run the everRun Availability Console.

everRun Enterprise ["System Hardware " on page 22](#) requirements are described below. See ["System Software" on page 23](#) for software requirements.

System Hardware

Supported Servers

Stratus everRun Enterprise software will run on any systems listed on the [Red Hat® Linux Hardware Catalog](#) that contains one of the following:

- One or two Intel® Xeon® E3-XXXX Processors or Intel Xeon E3-XXXX v2 Processors
- One or two Intel Xeon E5-XXXX Processors or Intel Xeon E5-XXXX v2 Processors

A second computer with identical processors is required for use as a redundant server for Protected Virtual Machines (or *PVMs*—virtual machines that are protected by Stratus everRun software). The CPUs for every host computer must have hardware support for virtualization enabled in the BIOS.

RAM

A minimum of 8 GB of RAM (physical memory) is recommended.

Disk Space Requirements

Only internal disks are supported. A minimum of two drives per physical machine is required for fault-tolerant operation.

50 GB is required for the host CentOS operating system and everRun software in the host domain including space for logs. Allow 10 GB minimum (boot disk) for each VM. Additional storage is needed for applications and data on each VM.

Network

The minimum network configuration includes two ports: one for A-link and one for a shared management/business link.

An optimal network configuration includes two 10-GbE network ports for A-Links (one of which also serves as `priv0`, the private network), a network interface for the management network, and as many

business/production ports as your protected VMs may need. If planning to run multiple protected VMs, consider adding pairs of A-Links, up to the supported total of four pairs.

All network components in a SplitSite configuration must have more than 155 Mbps minimum capacity, end-to-end. When fault-tolerant SMP is used, the A-Link networks must be a minimum of 1Gbps.

See ["Network Architecture Overview" on page 16](#), ["A-Link and Private Networks" on page 16](#), and ["Business and Management Networks" on page 17](#) for more information.

IP Addresses

Each everRun Enterprise host 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" on page 46](#) for more information.

Ports

everRun systems use port 443 in the local firewall for HTTPS communications, port 22 for ssh, and 5900 for VNC with Linux VMs. Firewalls must allow traffic through the appropriate ports. Firewalls must permit everRun-protected VMs to contact quorum service computers using UDP port 4557.

System Software

See ["Compatible Guest Operating Systems" on page 316](#)

Related Topics

["Physical Machine System Requirements" on page 317](#)

["Important Physical Machine and Virtual Machine Considerations" on page 319](#)

["Virtual Machine Recommendations and Limits" on page 319](#)

["Planning Virtual Machine Resources" on page 120](#)

["Configuring IP Settings" on page 64](#)

Storage Requirements

An everRun Enterprise system has the following storage requirements and recommendations:

- Each physical machine must contain at least two physical disks.
- Stratus strongly recommends that your system uses a storage RAID controller.
 - If the system has a single logical disk, Stratus strongly recommends that you configure the RAID controller so that logical disks presented to the host are backed by redundant physical drives.
 - Stratus strongly recommends that RAID controllers have a battery-backed write cache.
 - You must configure the RAID controller to boot off the first logical disk.

After confirming that your storage configuration meets these requirements, return to ["Site and System Preparation" on page 32](#).

Related Topics

["everRun Enterprise Storage Architecture" on page 15](#)

Memory Requirements

A minimum of 8 GB of RAM (physical memory) is recommended. The total amount of memory available on an everRun system is equal to the minimum of the amount of memory presented by either physical machine (PM) in the system. For example, in a system where one PM has 32 GB memory and another PM has 16 GB memory, the total amount of memory is 16 GB (least memory of either PM).

Related Topics

["Planning Virtual Machine Memory" on page 122](#)

General Network Requirements and Configurations

This topic discusses general network requirements and provides some recommended network configurations.

Requirements

Before you install everRun Enterprise software, make sure your network meets the following requirement:

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

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

- ["A-Link and Private Network Requirements" on page 26](#)
- ["Business and Management Network Requirements" on page 25](#)
- ["SplitSite Network Requirements" on page 27](#)

Recommended Configurations

Recommendations for possible configurations follow:

- If your system has **two 1 Gb** and **two 10 Gb** Ethernet ports:
 - Set one 10 Gb port as the private network (priv0).
 - Set the other 10 Gb port as an A-Link network.
 - Set one 1 Gb port as the management link.
 - Set the other 1 Gb port as a business link.

If your system has **four Ethernet ports of the same type** (for example, four 1 Gb or four 10 Gb interfaces):

- Set one port as the private network (priv0).
- Set one port as an A-Link network.
- Set one port as the management link.
- Set one port as a business link.



Note: A system with four 1 Gb Ethernet ports may not provide sufficient throughput for acceptable performance. The system may require 10 Gb add-on cards to achieve acceptable performance.

Business and Management Network Requirements

Business and management networks have the following requirements:

- Uses IPv6 link-local addressing.
- The speed of business or management networks should be less than or equal to the speed of A-Link networks.
- No support for bonding or VLAN trunking.
- 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 everRun Availability Console, use biz0:0, which is the IPv4 address that migrates to the primary management PM. Each PM also has its own IPv4 address (ibiz0) on the management network.
- Each PM requires at least one business network (specifically, the management network), with a maximum of 20 business networks.

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. An everRun 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 everRun 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 Protected Virtual Machine (PVM) to contain the following lines, save the file, and reboot the PVM.
 - `net.ipv4.conf.default.rp_filter = 2`
 - `net.ipv4.conf.all.rp_filter = 2`
- 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 everRun 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 everRun PM to the other, the VM may experience a blackout in which network traffic is not properly directed to and from the VM.

A-Link and Private Network Requirements

A-Link and private networks have the following requirements:

- Uses IPv6 link-local addressing.
- All A-Link and private networks on one PM of an everRun system must be in the same L2 broadcast domain as its matching links on the other PM, without any protocol filtering.
- Ethernet packets sent between two everRun PMs must not be obstructed or rate-limited. Ensure that they are not routed or switched by any L3 network infrastructure.
- One to eight A-Link networks per PM; however, a minimum of two is recommended.
- Uses 1 Gb to 10 Gb Ethernet ports. 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. A-Link networks are not required to be directly connected; instead, they can connect to a network switch.
- Private networks have no network hosts connected other than the everRun end-points.
- The system assigns each VM to a minimum of one or a maximum of two A-Link networks. However, each A-Link network can have multiple VMs assigned to it.

Related Topics

["A-Link and Private Networks" on page 16](#)

SplitSite Network Requirements

This topic describes the requirements for networks in a SplitSite configuration.

- ["A-Link Network Requirements" on page 27](#)
- ["Private Network Requirements" on page 28](#)
- ["Business Network Requirements" on page 28](#)
- ["Management Network Requirements" on page 29](#)

A-Link Network Requirements

A-Link networks in a SplitSite configuration require the following:

- NICs must be at least 1 Gb and full-duplex; use 10 Gb, if possible.
- For systems running FT-protected virtual machines (VMs), A-Links require:
 - A minimum bandwidth of 1 Gbps per VM
 - A minimum inter-site latency of 2 ms, round-trip time
- For systems running only HA-protected VMs, A-Links require:
 - A minimum bandwidth of 155 Mbps per VM
 - A minimum inter-site latency of 10 ms, round-trip time
- Do not use a common card (multiport NIC) for both A-Links.
- A-Links can be dedicated point-to-point fiber connections. If they are not, they must be configured on a VLAN. The A-Links can share a single VLAN, or they can use separate VLANs. Multiple everRun Enterprise systems can share the same VLAN(s) for use by the A-Links.

Private Network Requirements

The private network in a SplitSite configuration requires the following:

- NICs must be at least 1 Gb and full-duplex; use 10 Gb, if possible.
- A minimum bandwidth of 155 Mbps per VM.
- A minimum inter-site latency of 10 ms, round-trip time. Switches and/or fiber-to-copper converters connected to the private network must be non-routed and non-blocking, with a round-trip latency that does not exceed 10 ms. Calculate latency at 1 ms for each 100 miles of fiber, and approximately 1 ms for each non-routed, non-blocking switch or fiber converter.
- The private network can be a dedicated point-to-point fiber connection. If it is not, it must be configured on a private VLAN. VLANs used to connect the private-network ports must not add any filtering on any of the network equipment between the two VLAN switch ports that are connected to the everRun Enterprise PMs.

Business Network Requirements

Business networks in a SplitSite configuration require the following:

- Configure the network on a business VLAN. The business network for both nodes must be on this VLAN.
- The nodes must be in the same layer-2 multicast domain.

- Connect the business networks on each PM to a switch separate from the other PM's switch.
- An everRun Enterprise system requires at least one business network. All of these requirements apply to each business network.

Management Network Requirements

A management network in a SplitSite configuration requires the following:

- 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.

Related Topics

["SplitSite Configurations" on page 13](#)

["Network Architecture Overview" on page 16](#)

everRun Enterprise Availability Console Requirements

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

- Your computer must be able to access the subnet containing the everRun system.
- Use a supported browser. See ["Compatible Internet Browsers" on page 29](#).
- Make sure your computer has a recent release of Java 7 or later. Your browser may require you to update to the latest version. Java downloads are available at <http://www.java.com>.

For more information, see ["Using the everRun Availability Console" on page 51](#)

Compatible Internet Browsers

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

The following browsers are compatible with everRun Enterprise systems.

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

Java™ Requirements

Your system must be running an up-to-date version of Java. If you are running an outdated version, you may see a warning when using a wizard or other feature of the everRun Availability Console. If you continue using the feature, it is likely to hang. The warning will advise you to install the latest version of Java and either:

- Reduce your Java security settings to medium
- Add your everRun Enterprise system to the Exception Site List
- Or add a certificate as a Signer CA in Java, using the [link in the message](#)

Quorum Servers Considerations

Enabling or configuring quorum service is a post-install configuration task. Stratus recommends configuring two quorum service computers: a preferred quorum server and an alternate. For an overview of quorum servers, see ["Quorum Servers" on page 14](#).

Quorum service software, if deployed, can be installed on any general-purpose computer or laptop running the Windows operating system with these requirements:

- OS: Windows XP (SP2 or newer), Windows Server 2003, Windows Vista, or Windows 7, or Windows Server 2008; always powered on
- Disk space: 100 MB (minimum)
- NIC: at least one (1)
- Connectivity: Available to the everRun configuration via the management network

To install quorum server software:

¹IE8 is not recommended and it does not support some everRun Enterprise functionality.

1. From the **Drivers and Tools** section of the **everRun Enterprise Downloads and Support** page at <http://www.stratus.com/go/support/everrun>, download the quorum server software installer file to the quorum server.
2. On the quorum server, double click the installer file.



Note: When upgrading to a more recent version of quorum server software, you do **not** need to uninstall the previous version.

In addition, consider the following for quorum service best practices:

- Configure two quorum service computers with minimum common networking between the quorum computers and the hosts.
- As installed, protected VMs contact quorum service computers using UDP port 4557. Firewalls must permit everRun-protected VMs to contact quorum service computers using UDP port 4557. (If this port assignment conflicts with your local infrastructure, you can use different port numbers when you configure the quorum servers using the everRun Availability Console.)
- Quorum service computers should not reside on the same site as any host when deploying in a SplitSite. If both preferred and alternate quorum computers fail for a common reason, VMs will gracefully downgrade redundancy, and then continue to operate using one host, pending recovery of a quorum computer. However, if a host and the elected quorum computer fail for a common reason, the VMs running on the surviving server shut themselves down. For additional information about quorum servers and SplitSite configurations, see ["SplitSite Network Requirements" on page 27](#) and ["SplitSite Configurations" on page 13](#).
- If the preferred and alternate quorum service computers must reside on a common site, power them from separate AC power sources (phases) or configure them on separate UPS devices.

Related Topics

["Physical Machines and Virtual Machines" on page 8](#)

["Configuring Quorum Servers" on page 67](#)

["The everRun Availability Console" on page 51](#)

["Using the everRun Availability Console" on page 51](#)

Power Requirements and Considerations

To ensure maximum availability, Stratus strongly recommends that everRun Enterprise's fault-tolerant (FT) software run on physical machines (PMs) that are powered by redundant power supplies. In addition, each PM power supply should connect to a separate power source.

See ["Connecting Power" on page 33](#) for illustrations of some sample power-connection configurations.

Also, refer to your server vendor documentation for other power-related information.

Software Installation

When you install the everRun Enterprise software for the first time:

1. Prepare your site and system for the installation. See ["Site and System Preparation" on page 32](#).
2. Connect power to your system. See ["Connecting Power" on page 33](#).
3. Install the everRun software. See ["Installing everRun Enterprise Software" on page 36](#).

When the software installation is complete, see ["Post-Installation Tasks" on page 45](#).

Related Topics

["Upgrading everRun Enterprise Software" on page 95](#).

Site and System Preparation

Before you install everRun Enterprise software, make sure that your site and system meet the following requirements:

- The system meets all of the requirements described in ["System Requirements Overview" on page 22](#).
- The storage configuration meets all of the requirements described in ["Storage Requirements" on page 23](#).
- Provide keyboard and console access to each physical machine. This access can be in the form of a physical keyboard and monitor, a keyboard-video-mouse switch, or a properly-configured remote-management card capable of providing remote console and keyboard access. Connect the keyboard/console access as described in the vendor's documentation (for example, through direct VGA or USB connections).



Note: You cannot install everRun software from a serial console.

- Provide a remote management computer for the everRun Availability Console, and make sure it meets all of the requirements described in ["everRun Enterprise Availability Console Requirements" on page 29](#).
- Determine the best configuration for your network. See ["General Network Requirements and Configurations" on page 24](#).
- Use either an internal DVD drive or a USB DVD drive for the installation.

After confirming that your site and system meet the preceding requirements, return to ["Software Installation" on page 32](#).

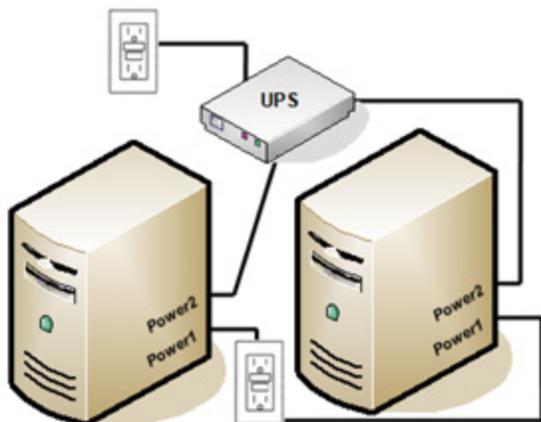
Connecting Power

To connect power, configure your everRun Enterprise server with redundant power supplies connected to separate sources. After connecting power, return to ["Software Installation" on page 32](#).

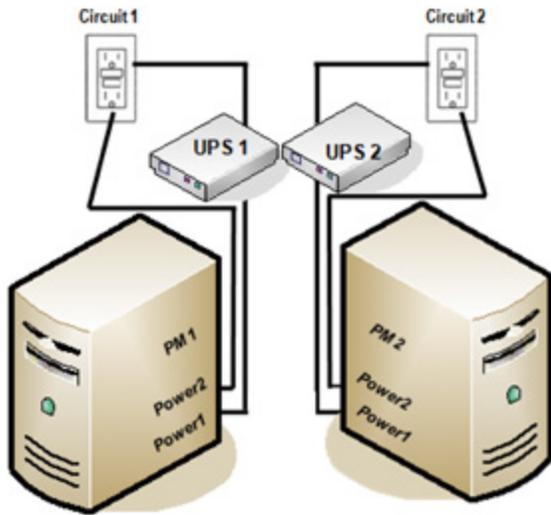
UPS (Optional)

The illustrations show how to connect one or two optional uninterruptible power supplies (UPS) to the everRun system.

Single UPS:



Dual UPS:



Related Topics

["Power Requirements and Considerations" on page 31](#)

Obtaining everRun Enterprise Software

Stratus provides the everRun Enterprise installation media as an ISO image. You can boot from it directly, or you can burn it to a DVD.



Note: You cannot boot the ISO image from a flash drive.

To obtain the ISO image:

1. From any computer connected to the Internet, go to the **everRun Enterprise Downloads and Support** page at <http://www.stratus.com/go/support/everrun>.
2. To download the everRun Enterprise software ISO image (**everRun_Enterprise_install-7.x.x.x-xx.iso**), under **Product Download**, click **everRun Enterprise 7.x.x.x ISO Image**. Save the ISO image.



Note: Depending on your Internet connection, the download can take up to 30 minutes.

Occasionally a file can be corrupted during the download process. To verify that the downloaded file is not corrupted, perform steps 3 through 7, otherwise skip to step 8.

3. To download the **MSUMS** file, under **Product Download**, click **everRun Enterprise 7.x.x.x md5sum**. Save the file.

4. From the Internet, find the **md5sum.exe** file and download it to the directory containing the everRun Enterprise ISO image.
5. If you have a Windows-based system, open a command prompt.
6. Go to the directory containing the **md5sum.exe**, **md5**, and ISO image files, and then type the following command to check the ISO image's status:
 - On Windows: **md5sum.exe -c everRun_Enterprise_install-7.x.x.x.xxx.md5**
 - On Linux: **md5sum -c everRun_Enterprise_install-7.x.x.x.xxx.md5**
7. If the command succeeds (that is, it returns the message `everRun_Enterprise_install-7.x.x.x.-xxx.iso: OK`), go to the next step. If the command fails, repeat the download.
8. When the validation is complete, do one of the following:
 - Burn the ISO image to a DVD, and then perform the next step in ["Installing everRun Enterprise Software" on page 36](#).
 - If you are not burning the ISO image to a DVD, perform the next step in ["Installing everRun Enterprise Software" on page 36](#).

Configuring the BIOS

Before you perform the software installation, you must modify certain BIOS settings. You can also modify some optional (but recommended) BIOS settings.

After you modify the BIOS settings, save them and perform the next step in the installation procedure (either ["Installing Software on the First PM" on page 39](#) or ["Installing Software on the Second PM" on page 43](#)



Note: This topic provides general information about BIOS settings. Because BIOS settings--including the setting names--vary, refer to the manufacturer's documentation for specific instructions for modifying the BIOS settings.

Required Settings

The following BIOS settings are required.

First Boot	Controls which device boots the operating system. Set the first boot device to the
------------	------------------------------------------------------------------------------------

Device	Optical Drive.
Vir- tualization Technology	Allows the processor to use virtualization technology. Set this to Enabled.
Execute- Disable Bit Capability	Allows the processor to classify areas in memory where application code can or cannot execute. Set this to Enabled to help prevent malicious code attacks.

Recommended Settings

The following BIOS settings are optional but recommended.

AC Power Recovery	Determines whether the server automatically powers on and boots after a power cycle. The recommended setting is ON.
F1/F2 Prompt on Error (Dell systems only)	Terminates booting if an error is detected during the process. Set this to Dis-able, as the everRun system may be able to provide more information once the server is running.

Installing everRun Enterprise Software

Follow these instructions to install everRun software for the first time on a system.



Warning: Installing everRun software erases all hard drives.

To install everRun software for the first time:

1. On a remote management computer, obtain the everRun Enterprise software. See "[Obtaining ever-Run Enterprise Software](#)" on page 34.
2. On your everRun system:
 - a. Provide keyboard and console access to your physical machines (PMs), if you have not already done so (see "[Site and System Preparation](#)" on page 32).

- b. Connect Ethernet cables for the networks you are configuring. See ["Connecting Ethernet Cables" on page 37](#).
3. Perform the installation on the first PM. See ["Installing Software on the First PM" on page 39](#).
4. After you have finished installing the software on the first PM, perform the installation on the second PM. See ["Installing Software on the Second PM" on page 43](#).

The installation is now complete. To perform the required post-installation configuration steps, see ["Post-Installation Tasks" on page 45](#).

Connecting Ethernet Cables

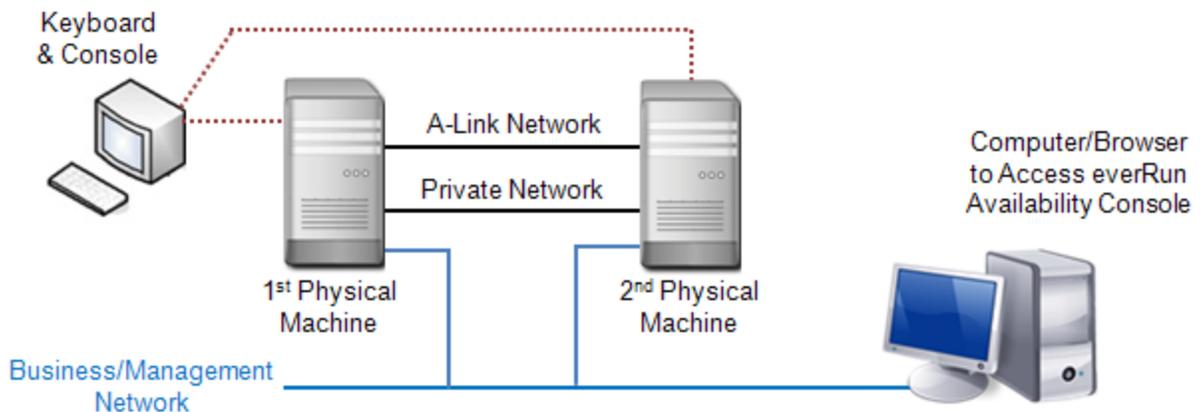
Before you install everRun Enterprise software for the first time, you need to connect Ethernet cables for your networks.



Note: To install additional networks after you have completed the software installation, see ["Connecting Additional Networks" on page 48](#).

On each physical machine (PM), assign one network port as the private network (priv0), and assign another network port as the management network (ibiz0). Although you can use any network port (1 Gb or 10 Gb) for the private network or management network, Stratus recommends that you use embedded network ports. Use CAT5E, CAT6, or CAT7 network cables for all network ports.

The following illustration shows an example of an everRun Enterprise network configuration.



Stratus recommends the following Ethernet cable configurations:

- For the private network, directly connect an Ethernet cable from any embedded port on the first PM to the same embedded port on the second PM. If you plan to use the private network as an A-Link, connect the cable to 10 Gb ports, if installed.
- For the management network, connect Ethernet cables from an embedded port on each PM to a network that is accessible from the remote management computer.



Note: Take note of the ports you used for the private and management networks. The installation software will prompt you for this information.

- For each A-Link network, connect an Ethernet cable from a port on the first PM to a port on the second PM, either directly or through a network switch.



Note: Stratus recommends that you configure at least one A-Link network in addition to the private network. See ["A-Link and Private Network Requirements" on page 26](#)

- For each business network, connect an Ethernet cable from a port on the first PM to a port on the second PM, through a network switch.

After you connect these Ethernet cables, perform the next step in ["Installing everRun Enterprise Software" on page 36](#).

Related Topics

["A-Link and Private Network Requirements" on page 26](#)

["Business and Management Network Requirements" on page 25](#)

["everRun Enterprise Availability Console Requirements" on page 29](#)

["Connecting Additional Networks" on page 48](#)

Installation Options

When you insert the everRun Enterprise DVD, the Welcome screen appears with the following list of installation-related options. Use the up and down arrow keys to select an option based on the task you want to perform. You can then press the **Tab** key to modify the command line. Finally, press the **Enter** key to boot the installation program from the DVD.

Task	Option	Description
Perform initial installation on the first PM	Install everRun, Create a new system	Deletes all partitions on all connected disks, installs the CentOS and everRun software, and creates a new system. See "Installing Software on the First PM" on page 39.
Perform initial installation on the second PM; replace a PM	Replace PM, Join system: Initialize data	Deletes all partitions on all connected disks, installs CentOS and everRun software, and attempts to connect to an existing system. See "Installing Software on the Second PM" on page 43 and "Replacing Physical Machines, Motherboards, NICs, or RAID Controllers" on page 182.
Recover failing PM	Recover PM, Join system: Preserving data	Preserves all data but re-creates the <code>/boot</code> and <code>root</code> file systems, reinstalls the CentOS and everRun software, and attempts to connect to an existing system. See "Recovering a Failed Physical Machine" on page 114.
Boot to rescue mode	Rescue the installed system	Boots to rescue mode.
Boot from local drive	Boot from local drive	Boots from a local drive.
Perform memory test	Memory test	Performs a memory test.

Installing Software on the First PM

This topic describes how to perform an initial installation of the everRun Enterprise software on node0, which is the first physical machine (PM).



Note: To perform an installation by mounting the ISO image, you must first configure your system's remote-management feature (for example, iDRAC on a Dell system). See the manufacturer's documentation for instructions.

To perform an initial installation of the software on the first PM:

1. Power on the first PM, if it is not already powered on, and either insert the installation software DVD or mount the ISO image.
2. As the system powers on, enter the BIOS and configure the required and optional BIOS settings as described in ["Configuring the BIOS" on page 35](#).
3. When the installation software loads, the **Welcome** screen appears and displays the options shown in ["Installation Options" on page 38](#). From this screen, you can choose to perform the initial installation in one of two ways:
 - **Method 1:** Installing via the user interface. This method is best for users who are not familiar with the installation process and who prefer to follow a GUI-based procedure with prompts.
 - **Method 2:** Installing via the command line. This method allows you to automate the installation. You can enter the IP settings in advance, and the installation proceeds without human intervention. This method is especially useful when you need to reinstall the software and you know all of the IP settings in advance.

Method 1: Installing via the User Interface

1. Use the arrow keys to select **Install everRun, Create a new system**, and press **Enter**.



Note: No action from you is required until the screen described in the next step appears.

2. The **Select interface for private Physical Machine connection** screen sets the physical interface to use for the private network. To use the first embedded port, use the arrow keys to select **em1** (if it is not already selected), and then press **F12** to save your selection and go to the next screen.

Notes:

1. If you are not sure of which port to use, use the arrow keys to select one of the ports, and click the **Identify** button. The LED on the selected port will then flash for 30 seconds, allowing you to identify it. Since the LED may also flash due to activity on that network, Stratus recommends that you leave the cable disconnected during the identification process. Reconnect the cable immediately after identification is complete.
2. If the system contains no embedded ports, select the first option interface instead.

3. The **Select interface for managing the system (ibiz0)** screen sets the physical interface to use for the management network. To use the second embedded port, use the arrow keys to select **em2** (if it is not already selected), and then press **F12** to save your selection and go to the next screen.



Note: If the system contains only one embedded port, select the first option interface. If the system contains no embedded ports, select the second option interface.

4. The **Select the method to configure ibiz0** screen sets the management network for node0 as either a dynamic or static IP configuration. Typically, you set this as a static IP configuration, so use the arrow keys to select **Manual configuration (Static Address)** and press **F12** to save your selection and go to the next screen. However, to set this as a dynamic IP configuration, select **Automatic configuration via DHCP** and press **F12** to save your selection and go to the next screen.
5. If you selected **Manual configuration (Static Address)** in the previous step, the **Configure em2** screen appears. Enter the following information and press **F12**.
 - IPv4 address
 - Netmask
 - Default gateway address
 - Domain name server address

See your network administrator for this information.



Note: If you enter invalid information, the screen redisplay until you enter valid information.

Method 2: Installing via the Command Line

1. Press the **Tab** key to bring up the command line.
2. Set the value for the private network (**priv0**).
 - To use the first embedded interface, type:
priv0=em1
 - To automatically select the default interface, type:
priv0=auto
 - To use the interface with a MAC address, type one of the following:
priv0=AA-BB-CC-DD-EE-FF or **priv0=AABBCCDDEEFF**
3. Set the value for the management network (**ibiz0**).
 - To use the second embedded interface with BOOTP:
ibiz0=em2:bootp
 - To automatically choose an interface and use DHCP:
ibiz0=auto:dhcp
 - To use a static configuration with IP address 10.83.51.116, netmask 255.255.0.0 , default gateway 10.83.0.1, and two DNS servers 134.111.24.254 and 134.111.18.14:
ibiz0=em2:10.83.51.116/16:10.83.0.1:134.111.24.254,134.111.18.14
 - To query the system administrator to configure the default interface:
ibiz0=auto
4. After entering the desired values on the command line, press **Enter**.
4. At this point, the installation continues without additional prompts. No action from you is required until the first PM reboots. After it reboots:

- a. Remove the DVD, or unmount the ISO image.
 - b. If you configured the IP address dynamically, record its IP address as described in ["Recording the Management IP Address" on page 43](#)
5. Perform the next step in ["Installing everRun Enterprise Software" on page 36](#).

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 you configured the management network 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 login screen similar to the following appears:

```

everRun Enterprise

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 ["Installing everRun Enterprise Software" on page 36](#) to see the next step.

Installing Software on the Second PM

This topic describes how to perform an initial installation of the everRun Enterprise software on node1, which is the second physical machine (PM).



Note: To perform an installation by mounting the ISO image, you must first configure your system's remote-management feature (for example, iDRAC on a Dell system). See the manufacturer's documentation for instructions.

To perform an initial installation of the software on the second PM:

1. Power on the second PM, if it is not already powered on, and either insert the installation software DVD or mount the ISO image.

2. As the system powers on, enter the BIOS and configure the required and optional BIOS settings as described in ["Configuring the BIOS" on page 35](#).
3. When the installation software loads, the **Welcome** screen appears and displays the options shown in ["Installation Options" on page 38](#). From this screen, you can perform the initial installation using either the user interface or the command line. This topic describes how to perform the installation with the user interface. To perform the installation with the command line, see "Method 2: Installing via the Command Line" in ["Installing Software on the First PM" on page 39](#).
4. Use the arrow keys to select **Replace PM, Join system: Initialize data**, and press **Enter**.



Note: No action from you is required until the screen described in the next step appears.

5. The **Select interface for private Physical Machine connection** screen sets the physical interface to use for the private network. To use the first embedded port, use the arrow keys to select **em1** (if it is not already selected), and then press **F12** to save your selection and go to the next screen.

Notes:



1. If you are not sure of which port to use, use the arrow keys to select one of the ports, and click the **Identify** button. The LED on the selected port will then flash for 30 seconds, allowing you to identify it. Since the LED may also flash due to activity on that network, Stratus recommends that you leave the cable disconnected during the identification process. Reconnect the cable immediately after identification is complete.
2. If the system contains no embedded ports, select the first option interface instead.

6. The **Select interface for managing the system (ibiz0)** screen sets the physical interface to use for the management network. To use the second embedded port, use the arrow keys to select **em2** (if it is not already selected), and then press **F12** to save your selection and go to the next screen.



Note: If the system contains only one embedded port, select the first option interface. If the system contains no embedded ports, select the second option interface.

7. The **Select the method to configure ibiz0** screen sets the management network for node1 as either a dynamic or static IP configuration. Typically, you set this as a static IP configuration, so use the arrow keys to select **Manual configuration (Static Address)** and press **F12** to save your selection and go to the next screen. However, to set this as a dynamic IP configuration, select **Automatic configuration via DHCP** and press **F12** to save your selection and go to the next screen.
8. If you selected **Manual configuration(Static Address)** in the previous step, the **Configure em2** screen appears. Enter the following information and press **F12**.
 - IPv4 address
 - Netmask
 - Default gateway address
 - Domain name server address

See your network administrator for this information.



Note: If you enter invalid information, the screen redisplay until you enter valid information.

9. At this point, the installation continues without additional prompts. No action from you is required until the second PM reboots. After it reboots:
 - a. Remove the DVD, or unmount the ISO image.
 - b. If you configured the IP address dynamically, record its IP address as described in ["Recording the Management IP Address" on page 43](#)
10. Perform the next step in ["Installing everRun Enterprise Software" on page 36](#).

Post-Installation Tasks

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

- ["Obtaining System IP Information" on page 46](#)
- ["Logging on to the everRun Availability Console for the First Time" on page 46](#)

- Configuring Required System Preference
 - ["Configuring Date and Time" on page 68](#)
 - ["Configuring Remote Support Settings" on page 77](#)
 - ["Configuring Quorum Servers" on page 67](#)
 - ["Specifying Owner Information" on page 59](#)
- ["Managing Local User Accounts" on page 92](#)
- ["Resolving Outstanding Alerts on the Dashboard" on page 54](#)
- ["Connecting Additional Networks" on page 48](#)

Obtaining System IP Information

After you install the everRun Enterprise software, you need the node0 IP address to log on to the everRun Availability Console for the first time (see ["Logging on to the everRun Availability Console for the First Time" on page 46](#)). 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 IP addresses (see ["Recording the Management IP Address" on page 43](#)), which helps the network administrator determine system IP information.

Obtain system IP values for the following:

- IP Address—The system IP address must be a static IP address. Do not use a dynamic IP address.
- NetMask
- Gateway Address
- Primary DNS
- Secondary DNS

Related Topics

["Software Installation" on page 32](#)

["Post-Installation Tasks" on page 45](#)

Logging on to the everRun Availability Console for the First Time

After completing the installation of the everRun Enterprise software, log on to the everRun Availability Console to accept the end-user license agreement (EULA) and to manage the everRun Enterprise system.

Prerequisites: To log on to the everRun Availability Console the first time, you need the following:



- The node0 (primary) IP address—You note this address during installation. See ["Recording the Management IP Address" on page 43](#).
- The system IP address—The network administrator provides this information. See ["Obtaining System IP Information" on page 46](#).
- The license **.KEY** file that you received from Stratus when you purchased the everRun Enterprise software—You must upload this file to the everRun Availability Console to complete the initial logon. Locate the file before you begin the initial logon.

To log on to the everRun Availability Console for the first time

1. From the remote management computer, type the IP address of node0 (primary) into a browser address bar.

The logon page of the everRun Availability Console appears.

2. Enter **admin** for the **Username** and **admin** for **Password**, and then click **LOGIN**

The Stratus everRun Enterprise EULA appears.

3. Read the EULA and then, if appropriate, click **Accept** to accept it.

The **INITIAL CONFIGURATION** page appears.

4. Under **NOTIFICATIONS**, the box for **Enable Support Notifications** is checked, by default. If you do not want the everRun Enterprise 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" on page 77](#)).

5. Under **SYSTEM IP**, enter the following values for the system IP configuration:

- IP Address (enter the address you obtained from your network administrator)
- NetMask
- Gateway Address
- Primary DNS
- Secondary DNS

After you have entered the network information, click **Continue**.

6. The **Portal Restart Required** window appears. After waiting a minute (as indicated in the window), click **OK** to refresh the console and to continue.
7. The **LICENSE INFORMATION** window appears. Under **Upload License Key**, click **Browse** and navigate to the license **.KEY** file that you received from Stratus. Select the license file and click **Upload**.

The initial logon is complete, and the everRun Availability Console appears. Bookmark or otherwise make note of the system IP address for use when logging in to the console in the future.

For security, change the default user login name and password for the **admin** account on the **Users & Groups** page. See ["Managing Local User Accounts" on page 92](#).

Related Topics

["Software Installation" on page 32](#)

["The everRun Availability Console" on page 51](#)

["Using the everRun Availability Console" on page 51](#)

Connecting Additional Networks

The everRun Enterprise installation software connects networks for all network ports that are physically connected at the time of the installation. This topic describes how to connect additional networks after the software installation is complete.

To connect a network:

1. Connect an Ethernet cable from a port on the first PM to a port on the second PM. Ideally, use the same NIC slot and port number in each PM. Connect the cable either directly (for an A-Link network) or through a network switch (for an A-Link or business network).
2. In the everRun Availability Console, go to the **Networks** page.
 - a. The new shared-network name should appear within a minute or so. If not, either your cable is on different subnets or the NIC ports between the PMs are incompatible (for example, if one end is connected to a 10 Gb port and the other end is connected to a 1 Gb port).
 - b. Click the **Config** button to select whether the network should be an A-Link or a business network. If the connection is direct, the network must be an A-Link. Otherwise, the network can

be either an A-Link or a business network.

- c. Verify that the new shared network displays a green check.
3. Connect additional network cables to both PMs, one pair at a time. Ideally, use the same NIC slot and port number in each PM.

Related Topics

["Connecting Ethernet Cables" on page 37](#)

["A-Link and Private Network Requirements" on page 26](#)

["Business and Management Network Requirements" on page 25](#)

["General Network Requirements and Configurations" on page 24](#)

3

Chapter 3: Using the everRun Availability Console

The everRun Availability Console is a browser-based interface that provides management and monitoring of an everRun Enterprise system from a remote management computer. For an overview of the console, see ["The everRun Availability Console" on page 51](#).

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

- ["The Dashboard Page" on page 53](#)
- ["The System Page" on page 55](#)
- ["The Preferences Page" on page 57](#)
- ["The Alerts Page" on page 79](#)
- ["The Audits Page" on page 80](#)
- ["The Physical Machines Page" on page 80](#)
- ["The Virtual Machines Page" on page 84](#)
- ["The Volumes Page" on page 88](#)
- ["The Storage Group Page" on page 89](#)
- ["The Networks Page" on page 90](#)
- ["The Virtual CDs Page" on page 90](#)
- ["The Upgrade Kits Page" on page 91](#)
- ["The Users & Groups Page" on page 91](#)

The everRun Availability Console

The everRun Availability Console is a browser-based interface that provides management and monitoring of an everRun Enterprise 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 everRun Availability Console, see ["everRun Enterprise Availability Console Requirements" on page 29](#).

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

- Read system alerts from the Dashboard. See ["The Dashboard Page" on page 53](#).
- View VM, CPU, memory, and storage statistics, and reboot or shutdown the system from the System page. See ["The System Page" on page 55](#).
- Set preferences for the system, diagnostics, notifications (e-Alerts and SNMP configuration), and remote support (notification and access). System preferences include owner information and configuration values for IP address, quorum services, date and time, etc. See ["The Preferences Page" on page 57](#).
- View alerts and audit logs. See ["The Alerts Page" on page 79](#) and ["The Audits Page" on page 80](#).
- Monitor, manage, and maintain resources:
 - PM status, storage, disks, network, and sensors: see ["The Physical Machines Page" on page 80](#).
 - VM status and management tasks such as creating, importing/restoring, managing, and maintaining VMs: see ["The Virtual Machines Page" on page 84](#).
 - Volumes, including their state, size, and storage group: see ["The Volumes Page" on page 88](#).
 - Storage groups, including name, size used, size, and number of volumes: see ["The Storage Group Page" on page 89](#).
 - Networks, including state, physical interface, speed, MAC address, and network bandwidth: see ["The Networks Page" on page 90](#).
 - Virtual CDs, including their state, name, size, and storage group: see ["The Virtual CDs Page" on page 90](#).

- Monitor and manage upgrade kits, and users and groups in the LIBRARY. See ["The Upgrade Kits Page" on page 91](#) and ["The Users & Groups Page" on page 91](#).

Related Topics

["Logging on to the everRun Availability Console for the First Time" on page 46](#)

["Logging on to the everRun Availability Console" on page 53](#)

["Using the everRun Availability Console" on page 51](#)

Logging on to the everRun Availability Console

Log on to the everRun Availability Console to manage the everRun Enterprise system. Using the console, you can manage the system, including its physical machines (PMs), virtual machines (VMs), storage, and networks. You can also generate statistics and view alerts and logs.

To log on to the everRun Availability Console

1. Type the everRun Enterprise 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 everRun system's static IP address, supplied during installation.

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 everRun Availability Console for the First Time" on page 46](#)

["The everRun Availability Console" on page 51](#)

["Using the everRun Availability Console" on page 51](#)

The Dashboard Page

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

To display additional information about outstanding alerts, click an entry in the list of alerts or click an alert symbol (for example, ) in the everRun system diagram. The information includes:

- The component associated with the issue (for example, the everRun 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" on page 54](#)).

Understanding the everRun 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" on page 80](#)

["The System Page" on page 55](#)

["The Virtual Machines Page" on page 84](#)

Resolving Outstanding Alerts on the Dashboard

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

To resolve outstanding alerts

On the everRun Availability 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" on page 53](#)

The System Page

The **System** page displays information about the everRun Enterprise system, and enables you to reboot or shut down the system. To open this page, click **System** in the left-hand navigation panel.

The **System** page displays resource allocations for the everRun system.

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

- ["Rebooting the System" on page 55](#)
- ["Shutting Down the System" on page 56](#)

You perform many other administrative tasks on the everRun system using the everRun Availability Console. For information, see ["The everRun Availability Console" on page 51](#).

To manage everRun system resources, see ["Configuring System Resources" on page 69](#).

Related Topics

["Using the everRun Availability Console" on page 51](#)

Rebooting the System

Reboot the everRun system using the everRun Availability Console to safely restart both PMs without incurring downtime for the VMs.



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



Note: You can reboot the system only if both PMs are running, healthy, and not in maintenance mode.



Prerequisite: Confirm that both PMs are running before rebooting.

To reboot the everRun system

1. Select **System** in the left-hand navigation panel.
2. Click the **Reboot** button.

Rebooting can take up to 15 minutes. You can observe the process in the everRun Availability Console as the system's PMs sequentially enter and then exit maintenance mode (for information on maintenance mode, see ["Maintenance Mode" on page 109](#)).

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 everRun Availability Console" on page 51](#)

["The System Page" on page 55](#)

["Using the everRun Availability Console" on page 51](#)

Shutting Down the System

Use the everRun Availability Console to shut down the everRun Enterprise 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 everRun system. Make sure both PMs are running before shutting down.

Cautions:



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

To shut down the everRun system

1. 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.

You can observe some of shutdown process in the everRun Availability Console as the system's PMs sequentially enter maintenance mode (for information on maintenance mode, see ["Maintenance Mode" on page 109](#)). When the system shuts down completely, though, the everRun Availability Console is unavailable and the masthead displays **Lost Communication**.

After the system shuts down, you lose the connection to the console. If the everRun system cannot shut down 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 everRun Availability Console. Click **Virtual Machines** in the left-hand navigation panel, select the VM, and then click **Power Off**.

Related Topics

["The everRun Availability Console" on page 51](#)

["The System Page" on page 55](#)

["Using the everRun Availability Console" on page 51](#)

The Preferences Page

The **Preferences** page enables you to configure everRun Enterprise 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 everRun system administrator. This information is also provided in response to Simple Network Management Protocol (SNMP) requests. See "Specifying Owner Information" on page 59 .

Preference	Description
IP Configuration	Allows you to view and specify the Internet Protocol (IP) address and network settings for the everRun system. See "Configuring IP Settings" on page 64 .
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 everRun Enterprise environment. See "Quorum Servers" on page 14 and "Configuring Quorum Servers" on page 67 .
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 everRun system. See "Configuring Date and Time" on page 68 .
System Resources	Allows you to specify the number of virtual CPUs (VCPUs) and the amount of memory reserved for the everRun Enterprise software. See "Configuring System Resources" on page 69 .
Import Settings	Allows you to enable encryption for import procedures in the everRun Availability Console. See "Configuring the Virtual Machine Import Option" on page 70 .
Diagnostics	
Diagnostics	Allows you to generate a diagnostic file for your authorized Stratus service representative. See "Managing Diagnostic Files" on page 71 .
Notification	
e-Alerts	Allows you to enable email alerts (e-Alerts) for system administrators. See "Configuring e-Alerts" on page 73 .
SNMP Configuration	Allows you to enable Simple Network Management Protocol (SNMP) requests and traps for remote system monitoring. See "Configuring SNMP

Preference	Description
	Settings" on page 75.
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 everRun Enterprise system can send notifications to your authorized Stratus service representative about problems with the system. See "Configuring Remote Support Settings" on page 77.
Proxy Configuration	Allows you to configure proxy settings for the everRun Enterprise system if your organization requires a proxy server to access the Internet and you have a service agreement with Stratus or another authorized everRun service representative. everRun software uses proxy server information for support notification messaging and remote support access features. See "Configuring Internet Proxy Settings" on page 78.

Related Topics["The everRun Availability Console" on page 51](#)["Using the everRun Availability Console" on page 51](#)**Specifying Owner Information**

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

This contact information is available in the everRun Availability 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**.

Managing the everRun Enterprise Product License

Manage the product license for the everRun Enterprise system by:

- Uploading a license .key file saved on a computer.
- Downloading an activated license .key file to a computer and then uploading it to the everRun Enterprise system.
- Activating, renewing, or checking the status of an existing license.

When you purchase an everRun Enterprise system, Stratus provides you with a license .key file (via email). Save the license .key file to a location on a computer (not your everRun Enterprise system) that you can access when you need to upload (and activate) the license to the everRun Enterprise system.

If you do not yet have a license, or if you need to upgrade or renew a license or support, you must contact everRun Enterprise Customer Support or your authorized Stratus service representative. See the **everRun Enterprise Downloads and Support** page at <http://www.stratus.com/go/support/everrun>.

Your license is automatically activated/renewed each time you upload a license .key file to an everRun Enterprise system that has Internet connectivity to the Stratus `alas.stratus.com` server via port 443 (https). The everRun Enterprise system also attempts to activate/renew your license every 24 hours. If your everRun Enterprise system does not have Internet connectivity, you can manually download an activated .key file to a computer and then upload it to the everRun Enterprise system.

To upload a new license .key file to an everRun Enterprise system with internet connectivity

After you have saved a license.key file to a computer, use this procedure to upload the license.key file to the everRun Enterprise system. The everRun Enterprise system must have internet connectivity.

1. In the everRun Availability Console, click **Preferences** in the left-hand navigation panel.
2. On the **Preferences** page, click **Product License**.
3. Click the **New License** bar to display its options.
4. Under **Upload New License Key**, click **Browse** and navigate to the location of the license .key file on your computer. Select the license .key file and click **Open**. Then click **Upload** to upload the file to the everRun Enterprise system. The everRun Enterprise system contacts the Stratus server to activate the license.

To license an everRun Enterprise system with no Internet connectivity (but is connected to a computer that has Internet connectivity)

If your everRun Enterprise system is not connected to the Internet but has private intranet connectivity to a computer that is connected to the Internet, perform the following steps to download an activated license and then upload it to an everRun Enterprise system.

1. In the everRun Availability Console, click **Preferences** in the left-hand navigation panel.
2. On the **Preferences** page, click **Product License**.
3. Click the **License Check and Activation** bar to display its options.
4. Under Step 1, **Download Activated License Key**, click **Activated License** to activate and download a license .key file to a computer (not the everRun Enterprise system).

The **Opening av_number_A.key** dialog box appears. In the dialog box, select **Save File** and select a location on the computer to save the downloaded .key file. (Depending on the browser, the default location for saving the file may be the Downloads folder.)

5. Under Step 2, **Upload Activated License Key**, click **Browse** and navigate to the license .key file that you saved in the previous step. Then, click **Upload** to upload it to the everRun Enterprise system.

To license an everRun Enterprise system with no Internet connectivity

If your everRun Enterprise system is not connected to the Internet and has no private intranet connectivity to a computer that is connected to the Internet, perform the following steps to obtain an activated license and then move it to an everRun Enterprise system.

For this procedure;

- You will need a USB flash drive and two computers (A and B) in addition to the everRun Enterprise system.
- Computer A has internet access and has no connection to the everRun Enterprise system.
- Computer B has access to the everRun Availability Console on the everRun Enterprise system but both of these computers are **not** connected to the internet.

On Computer B

1. Insert a USB flash drive into a USB port.
2. Log on to the everRun Availability Console.

3. Click **Preferences** in the left-hand navigation panel.
4. On the **Preferences** page, click **Product License**.
5. Click the **License Check and Activation** bar to display its options.
6. Under Step 1, right click the **Activate License** link and select your browser's option to copy the link (for example, **Copy Link Location** or **Copy Link Address**).
7. Open a text editor (notepad.exe), paste the copied URL in to it, and save the contents of the editor to a text file on the USB flash drive.
8. Remove the USB flash drive.

On Computer A

1. Insert the USB flash drive into a USB port.
2. In a text editor, open the text file on the USB flash drive. Copy the URL in the text editor to the clipboard.
3. Open a web browser and paste the URL into the address bar. Press **Enter**. A license .key file will be downloaded.
4. Copy the license .key file to the USB flash drive.
5. Remove the USB flash drive.

On Computer B

1. Insert the USB flash drive into a USB port.
2. Click **Preferences** in the left-hand navigation panel.
3. On the Preferences page, click **Product License**.
4. Click the **License Check and Activation** bar to display its options.
5. Click **Browse**, navigate to the license .key file on the USB flash drive, and select it. Click **Open**.
6. In the **Product License** pane, click **Upload**.

To check the status of a license

Use this procedure to check the status of a license .key file that you have already uploaded on a computer that has Internet connectivity to the Stratus `alas.stratus.com` server via port 443 (https).

1. In the everRun Availability Console, click **Preferences** in the left-hand navigation panel.
2. On the **Preferences** page, click **Product License**.

Click the **License Check and Activation** bar.

3. Click **Check License Now**. The console displays the status of the license:

STATUS: License is activated and expires in *nn* days *nn* hours

TYPE OF LICENSE: Enterprise Edition (volume)

EXPIRATION: *month dd, 20yy, time*

LAST CHECK: *month dd, 20yy, time*

Asset ID: *asset_id*

License Activation Error Codes

If a license activation fails, the License Activation Server (or ALAS) returns one of the following numeric error codes.

2.1: ALAS_UNKNOWN_SITEID

The specified Asset ID key does not exist in the Stratus customer database Atlas. If the license was just created (for example, with trial IDs), the license information might not yet have propagated to ALAS. Wait 15 minutes and try again. If the activation fails again, contact your authorized Stratus service representative and provide them with the return code.

3.1: ALAS_INVALID_ARG

The ALAS URL was called without an Asset ID parameter. This error can occur with an improperly formed license key that does not include the Asset ID.

3.2: ALAS_INVALID_SITEID

The Asset ID parameter has been specified but does not contain a value. This error can occur with an improperly formed license key that includes a blank Asset ID.

3.3: ALAS_NO_SIGN

ALAS cannot communicate with the SSL certificate signing server.

3.4: ALAS_NO_ATLAS_UPDATE

ALAS failed to update activation information, OS release number, and/or other information in Atlas. This error occurs on the ALAS side of the license activation.

3.5: ALAS_NO_MORE_ACTIVATION

The site has exceeded the number of activations allowed (typically 3). If necessary, your authorized Stratus service representative can change the limit.

9.0: ALAS_UNKNOWN

Unknown error.

Related Topics

["The everRun Availability Console" on page 51](#)

["The Preferences Page" on page 57](#)

["Using the everRun Availability Console" on page 51](#)

Configuring IP Settings

Configure Internet Protocol (IP) settings for the everRun Enterprise 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.

During installation and post-installation of everRun Enterprise software, you configure three IP addresses: one for the everRun Enterprise system and one for each node. You can change the IP addresses and other IP settings after installation using the procedures below. You must specify a static IPv4 address for the everRun Enterprise system.



Note: The procedure you use to configure IP settings depends on whether the everRun Enterprise system stays in the same subnet or moves to a new subnet. Follow the appropriate procedure for your needs.

To change the system and/or node IP settings (same subnet)

The everRun Enterprise system and all virtual machines (VMs) continue to run throughout this procedure; however, the everRun Availability Console briefly loses its connection to the system if you change the system IP address. You can access the everRun Availability 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. Click the tab for one of the following: **System IP**, **Node0 IP**, or **Node1 IP**.
4. Type appropriate values for the following settings, if available on the tab:
 - **IP address** (within the same subnet)
 - **Gateway Address**
 - **Primary DNS**
 - **Secondary DNS**
 - **NetMask**
5. Click **Save** to save the values (or click **Reset** to restore previous values).

If you have changed the system IP address, the **Portal Restart Required** dialog box appears.

Wait about a minute and then click **OK**. This redirects the browser to the new system IP address.

6. If appropriate, click the next tab, repeating steps 4 and 5.

To change the system and/or node IP settings (new subnet)

Changing the system 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.



Caution: Perform this procedure only during a planned maintenance period since you must shut down the nodes.



Prerequisites: Obtain the following IP addresses (in IPv4 format) from a network administrator:

1. The existing IP address of the primary node—After you change the system IP address, you lose the connection to the everRun Availability Console. You use the existing IP address of the primary node to reconnect to the everRun Availability Console.
2. The new system IP address and, if applicable, other new system IP settings.
3. The new IP addresses for the nodes and, if applicable, other new node IP settings.

1. Click **Preferences** in the left-hand navigation panel, to open the **Preferences** page.
2. On the **Preferences** page, click **IP Configuration**.

3. Click the **System IP** tab.
4. Type the new system IP address (new subnet) for **IP address**. If appropriate, enter values for:
 - **NetMask**
 - **Gateway Address**
 - **Primary DNS**
 - **Secondary DNS**
5. Click **Save** to save the values.

The **Portal Restart Required** dialog box appears because you changed the system IP address. Click **OK**.



Note: When you click **OK**, you are redirected to the new system IP address, and your connection to everRun Availability Console is broken.

6. Open a web browser and enter the existing IP address of the primary node, to re-connect to the everRun Availability Console.
7. Shut down the system. See ["Shutting Down the System" on page 56](#).
8. Disconnect network cables, move the PMs to a new physical location, and/or attach the network cables to new switch ports, as appropriate for your needs.
9. Power on each PM by physically pushing the power button on each machine.
10. When the everRun Availability Console becomes available (within 10 to 15 minutes), open a web browser, type the new everRun system IP address and then log on.
11. Navigate to **IP Configuration** of **Preferences** (see steps 1 and 2, above).
12. Click the tab for **Node0** or **Node1**.
13. Type the new **IP address** for the node, and if appropriate, enter values for:
 - **Gateway Address**
 - **Primary DNS**
 - **Secondary DNS**
 - **Netmask**

14. Click **Save** to save the values.
15. Repeat steps 12, 13, and 14 for the other node.

Related Topics

["Software Installation" on page 32](#)

["The Preferences Page" on page 57](#)

["The everRun Availability Console" on page 51](#)

["Using the everRun Availability Console" on page 51](#)

Configuring Quorum Servers

When you log on to the everRun Enterprise system for the first time, configure quorum servers.



Prerequisite: Before you configure quorum servers, read ["Quorum Servers" on page 14](#) and ["Quorum Servers Considerations" on page 30](#).



Note: 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" on page 156](#) and ["Starting a Virtual Machine" on page 156](#).

To configure quorum servers

1. Log on to the everRun Availability Console.
2. Click **Preferences** in the left-hand navigation panel.
3. Click **Quorum Servers**.
4. Click **Add Quorum Server**.
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 value is 4557)—Type the port number if it is different from the default.

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**.

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 everRun Availability 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 everRun Availability Console" on page 51](#)

["The Preferences Page" on page 57](#)

["Using the everRun Availability Console" on page 51](#)

Configuring Date and Time

When you log on to the everRun Enterprise 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 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 everRun system.
- Configure all VMs to use the same NTP servers as those configured for the everRun system.

To configure date and time settings

1. Log on to the everRun system.
2. Click **Preferences** in the left-hand navigation panel.
3. On the **Preferences** page, click **Date & Time**.
4. In the **Date & Time** display, select a value for **Configure Time Zone** from its pull-down menu:
 - **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 everRun 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 everRun Availability Consolemasthead telling you that the system will reboot. In this case, the primary physical machine (PM) reboots and the secondary PM shuts down. While the primary PM reboots, you lose your connection to the everRun Availability 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 everRun Availability Console" on page 51](#)

["The Preferences Page" on page 57](#)

["Using the everRun Availability Console" on page 51](#)

Configuring System Resources

Configure system resources to specify how the everRun Enterprise system manages virtual CPUs (VCPUs) and memory. Use default values; change a value only if your service representative instructs you to.

To configure system resources for the everRun Enterprise system

1. In the everRun Availability Console, click **Preferences** in the left-hand navigation panel, to open the **Preference** page.
2. Click **System Resources**.
3. Modify the settings only if your service representative instructs you to:
 - **System VCPUs**, which sets the number of VCPUs reserved for the everRun Enterprise software. Values are **2** (the default) and **4**.
 - **System Memory**, which sets the amount of memory reserved for the everRun Enterprise software. Values are **1024 MB**, **2048 MB** (the default), and **4096 MB**
4. Scroll to the bottom of the **System Resources** section and click **Save** (or click **Reset** to restore the previously-saved values).

Related Topics

["The everRun Availability Console" on page 51](#)

["The Preferences Page" on page 57](#)

["Using the everRun Availability Console" on page 51](#)

Configuring the Virtual Machine Import Option

Configure the virtual machine import/export option to enable encryption that improves security for the everRun Enterprise system.

To configure an import option for the everRun system

1. In the everRun Availability Console, click **Preferences** in the left panel, to open the **Preferences** page.
2. Click **Import**.
3. Select a setting, as appropriate for your system:

- **Import** , which enables encrypted data communication using a secure version of the Hyper Text Transfer Protocol (HTTPS). Encryption can be time consuming, so enable it only when security is a concern. By default, this setting is disabled.

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

Related Topics

["The everRun Availability Console" on page 51](#)

["The Preferences Page" on page 57](#)

["Using the everRun Availability Console" on page 51](#)

Managing Diagnostic Files

Diagnostic files provide a snapshot of an everRun Enterprise system's log files and configuration information. This information enables your authorized Stratus service representative to resolve an issue with the system.

When you create diagnostic files, you can choose to include log files from the last 24 hours, the previous seven days, or all available log information and statistics for the everRun system. You can also choose to include only performance statistics.

For additional information, see:

- ["Creating a Diagnostic File" on page 71](#)
- ["Deleting a Diagnostic File" on page 73](#)
- ["Uploading a Diagnostic File to Customer Support" on page 72](#)

Related Topics

["The everRun Availability Console" on page 51](#)

["Using the everRun Availability Console" on page 51](#)

["The Preferences Page" on page 57](#)

Creating a Diagnostic File

Diagnostic files provide a snapshot of an everRun Enterprise system's log files and configuration information. You create a diagnostic file to help your authorized Stratus service representative resolve issues with the system.

To create diagnostic files

1. Click **Preferences** in the left-hand navigation panel, to open the **Preference** page.
2. In the **Diagnostics** category, click **Diagnostics**
3. Select an option from the pulldown menu:
 - **Minimal** diagnostics contain log information for the last 24 hours.
 - **Medium** diagnostics contain log information for the last 7 days.
 - **Full** diagnostics contain all available log information with statistics for the everRun system.
 - **Statistics** contain performance statistics for the last 7 days.
4. Click **Generate Diagnostic File**.
5. Upload the file to your authorized Stratus service representative, as described in ["Uploading a Diagnostic File to Customer Support" on page 72](#).

Related Topics

["The everRun Availability Console" on page 51](#)

["The Preferences Page" on page 57](#)

["Using the everRun Availability Console" on page 51](#)

Uploading a Diagnostic File to Customer Support

Upload a diagnostic file to the Stratus everRun Customer Support web site to help resolve an issue with the everRun Enterprise system. (To create a diagnostic file, see ["Creating a Diagnostic File" on page 71](#).)

To upload a diagnostic file to Customer Support

1. Click **Preferences** in the left-hand navigation panel, to open the **Preference** page.
2. In the **Diagnostics** category, click **Diagnostics**
3. Do one of the following:
 - If the everRun Enterprise system has Internet connectivity, upload the diagnostic file directly to the Stratus everRun Customer Support web site by clicking **Upload**.
 - If the everRun Enterprise system does not have Internet connectivity or if the **Upload** fails, you can manually upload the diagnostic file to the **Stratus Diagnostic File Upload** web page. First, click **Download** on the everRun Availability 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 File Upload** page, click **Browse**, select the file on the computer, and then click **Submit**.

If you need help with this procedure, call everRun Enterprise Customer Support at the phone number listed on the **everRun Enterprise Downloads and Support** page at <http://www.stratus.com/go/support/everrun>.

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 everRun Enterprise system, as described in ["Deleting a Diagnostic File" on page 73](#).

Related Topics

["The everRun Availability Console" on page 51](#)

["The Preferences Page" on page 57](#)

["Using the everRun Availability Console" on page 51](#)

Deleting a Diagnostic File

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

To delete a diagnostic file

1. Click **Preferences** in the left-hand navigation panel, to open the **Preference** page.
2. In the **Diagnostics** category, click **Diagnostics**
3. Select the diagnostic file and click **Delete**.

Related Topics

["The everRun Availability Console" on page 51](#)

["The Preferences Page" on page 57](#)

["Using the everRun Availability Console" on page 51](#)

Configuring e-Alerts

Configure email alerts (e-Alerts) to enable the everRun Enterprise 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.
 - **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 everRun 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.
 - **Connect using TLS**—Click this box if the SMTP server requires Transport Layer Security (TLS).
 - **Enable Authentication**—Click this box if the SMTP server requires authentication to send email and 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 everRun software generates a test alert that triggers the delivery of an e-Alert. Watch the Alerts History log (see ["The Alerts Page" on page 79](#)) for delivery status. A sample email with subject "Test Alert" is sent to all the email recipients.

You can also test e-Alerts by putting the secondary physical machine into maintenance mode (see ["Maintenance Mode" on page 109](#)), and then removing it from maintenance mode. Verify that you receive e-Alerts for both maintenance mode events.

Related Topics

["The everRun Availability Console" on page 51](#)

["The Preferences Page" on page 57](#)

["Using the everRun Availability Console" on page 51](#)

Configuring SNMP Settings

Configure Simple Network Management Protocol (SNMP) settings for your everRun Enterprise 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 everRun system to retrieve the values of objects listed in the Management Information Bases (MIBs) supported by the everRun software. These MIBs include an everRun-specific MIB that is a collection of objects describing the everRun system. For details about the everRun Enterprise MIB, see ["MIB File Contents" on page 324](#).
- **SNMP trap**—a message initiated by the everRun system after an event such as an alert that is then sent to an identified list of recipients, typically a network management station (NMS).

To specify the desired security parameters, you must edit the standard `/etc/snmp/snmpd.conf` file on both nodes. For example, to allow SNMP requests by any user using the default `public` community, comment out or delete the following lines from that file on each node:

```
com2sec notConfigUser default public
group notConfigGroup v1 notConfigUser
group notConfigGroup v2c notConfigUser
view systemview included .1.3.6.1.2.1.1
view systemview included .1.3.6.1.2.1.25.1.1
access notConfigGroup "" any noauth exact systemview none none
```

After you save the edited files, you must restart the `snmpd` process on each node by entering the following command:

```
service snmpd restart
```

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

To enable SNMP traps

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**, type the IP address or host name for each recipient, one per line.
6. Click **Save**. (Or click **Reset** to restore the previously saved values.)
7. Configure your organization's firewall to allow SNMP operations, as described below.
8. Generate a test alert, as described below.



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

To configure your firewall to allow SNMP operations

To enable SNMP management systems to receive alerts from and send traps to the everRun system, configure your organization's firewall to open the following ports:

Message Type: SNMP

Protocol: SNMP

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

To generate a test alert

Click **Generate Test Alert**. A test alert gets generated that triggers the delivery of SNMP traps. Watch the Alerts History log (see ["The Alerts Page" on page 79](#)) for delivery status. A sample SNMP trap is sent to all the recipients.

Configuring Remote Support Settings

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

To configure support configuration settings

1. In the everRun Availability Console, click **Preferences** in the left-hand navigation panel, to open the **Preference** page.
2. Under **Remote Support**, click **Support Configuration**.
3. Modify the settings as needed. See the descriptions below.
4. Click **Save** (or click **Reset** to restore the previously saved values).
5. Configure your organization's firewall to allow support messages, as described below.
6. Generate a test alert, as described below.



Note: When you enable or modify support configuration settings, generate a test alert to confirm that your authorized Stratus service representative can receive system health messages from your system.

Set values for the following settings, as appropriate for your system:

- **Enable Remote Support Access** allows your authorized Stratus service representative to remotely connect to the everRun system for troubleshooting purposes. Note that you can enable this setting and then disable the setting as needed.
- **Enable Notifications** allows the everRun 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.
 - **Enable Periodic Reporting** sends a daily summary of system information to help improve product and service quality.

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

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

Message Type: SNMP

Protocol: SNMP

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

To generate a test alert

Click **Generate Test Alert**. A test alert is generated which sends a support notification message.

Watch the Alerts History log for delivery status. A subsequent alert will be generated if the support notification fails.

Related Topics

["The everRun Availability Console" on page 51](#)

["The Preferences Page" on page 57](#)

["Using the everRun Availability Console" on page 51](#)

Configuring Internet Proxy Settings

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

A proxy server provides a secure bridge between the everRun Enterprise system and the Internet. everRun 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 everRun Availability Console" on page 51](#)

["The Preferences Page" on page 57](#)

["Using the everRun Availability Console" on page 51](#)

The Alerts Page

The **Alerts** page displays messages about events on the everRun Enterprise system.

To open the **Alerts** page, click **Alerts** in the left-hand navigation panel of the everRun Availability Console. (To view a log of user activity on the everRun system, see ["The Audits Page" on page 80.](#))

To view alert information, scroll through the alerts, which are, by default, listed in reverse chronological order. Click an alert to display information about the problem and resolution (if available), and whether **Support Notifications**, an **e-Alert**, or an **SNMP Trap** was sent for this alert.

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



- ["Configuring Remote Support Settings" on page 77](#)
- ["Configuring e-Alerts" on page 73](#)
- ["Configuring SNMP Settings" on page 75](#)

Related Topics

["The everRun Availability Console" on page 51](#)

["Using the everRun Availability Console" on page 51](#)

The Audits Page

The **Audits** page displays a log of user activity in the everRun Availability Console. To open this page, click **Audits** in the left-hand navigation panel. (To display information about events on the everRun system, see ["The Alerts Page" on page 79.](#))

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 everRun Availability Console was running.
- **Action**—The action performed in the everRun Availability Console.

Related Topics

["The everRun Availability Console" on page 51](#)

["Using the everRun Availability Console" on page 51](#)

The Physical Machines Page

The **Physical Machines** page enables you to manage the physical machines (PMs) in the everRun Enterprise 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** under **Name**. To interpret PM states and activities, see ["Physical Machine States and Activities" on page 83](#).

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

- **Action buttons:** Various action buttons appear, depending upon the state of the selected node. The **Work On** button () appears initially. To perform most maintenance tasks, click **Work On**, which places a node into maintenance mode (for information, see ["Maintenance Mode" on page 109](#)). To learn about additional PM actions available in maintenance mode, see ["Physical Machine Actions" on page 81](#) 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 the model, 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, controller and current action (if any) of storage.
 - **Network**, which displays the state, name, speed, and MAC address of networks.
 - **Sensors**, which displays the name and current state of sensors.
 - **Virtual Machines**, which displays the state, activity, and name of virtual machines.
 - **Details**, which displays the manufacturer, model, and serial number of the selected node.

Related Topics

["The everRun Availability Console" on page 51](#)

["Using the everRun Availability Console" on page 51](#)

Physical Machine Actions

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



Caution: Use the **Physical Machines** page of the everRun Availability Console when you perform maintenance on a PM. Do not use controls on the computer (for example, the power switch on the PC) because the everRun Availability Console protects the everRun 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 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" on page 109 .
The following actions are available 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" on page 109 .
 Shutdown	Shuts down a PM. The PM transitions to off (in Maintenance) . See "Shutting Down a Physical Machine" on page 112 .
 Reboot	Reboots the PM. The PM transitions to preparing for reboot (in Maintenance) . See "Rebooting a Physical Machine" on page 111 .
 Replace	Causes the everRun Enterprise software to delete the PM from the everRun system's database, so that you can replace the PM or one of its components. See "Replacing Physical Machines, Motherboards, NICs, or RAID Controllers" on page 182 .
The following action may be available when everRun software has removed a PM from service and powered it off, due to an excessive failure rate.	

Commands	Description
 Reset Device	Resets the mean time between failures (MTBF) counter for a PM so it can be brought back into service. See "Resetting MTBF for a Failed Physical Machine" on page 116 .

Related Topics

["The everRun Availability Console" on page 51](#)

["Using the everRun Availability Console" on page 51](#)

["The Physical Machines Page" on page 80](#)

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
	 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 Reset Device	everRun has powered off the PM because of an excessive failure rate. The PM is left powered off until Reset Device is clicked. See "Resetting MTBF for a Failed Physical Machine" on page 116 .
	 Booting	Finalize	PM is booting.
	 Rebooting	Finalize	PM is rebooting.

	 Running	Finalize Shutdown Reboot Recover Replace	PM is running in Maintenance Mode. See "Maintenance Mode" on page 109 .
-----------------------------------------------------------------------------------	-------------------------------------------------------------------------------------------	-------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------

Related Topics

["The everRun Availability Console" on page 51](#)

["Using the everRun Availability Console" on page 51](#)

["The Physical Machines Page" on page 80](#)

The Virtual Machines Page

Use the **Virtual Machines** page to manage the virtual machines (VMs) running on your everRun Enterprise system. To open this page, click **Virtual Machines** in the left-hand navigation panel of the everRun Availability 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" on page 87](#). To learn more about the controls on this page, see ["Virtual Machine Actions" on page 85](#) or the help topic for a specific task.

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

- ["Creating and Migrating Virtual Machines" on page 125](#)
- ["Opening a Virtual Machine Console Session" on page 158](#)
- ["Reprovisioning Virtual Machine Resources" on page 161](#)
- Controlling the power state of a VM, as described in:
 - ["Starting a Virtual Machine" on page 156](#)
 - ["Shutting Down a Virtual Machine" on page 156](#)
 - ["Powering Off a Virtual Machine" on page 157](#)
- ["Removing a Virtual Machine" on page 160](#) or ["Renaming a Virtual Machine" on page 159](#)

- Managing advanced tasks or troubleshooting, as summarized in ["Advanced Topics \(Virtual Machines\)" on page 173](#)
- Viewing information about a VM, including its name, description, and resources in the tabs of the bottom pane

Related Topics

["Managing Virtual Machines" on page 119](#)

["Using the everRun Availability Console" on page 51](#)

Virtual Machine Actions

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

Action	Description
 Create	Launches the Create VM Wizard. See "Creating a New Virtual Machine" on page 125 .
 Import/Restore	Imports or restores a VM from a set of OVF and VHD files. See "Creating and Migrating Virtual Machines" on page 125 . 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.	
 Console	Opens a console for the selected VM. See "Opening a Virtual Machine Console Session" on page 158 .
 Shutdown	Shuts down the selected VM. See "Shutting Down a Virtual Machine" on page 156 .

Action	Description
 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 shut-down. See "Powering Off a Virtual Machine" on page 157 .
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" on page 161
 Start	Boots the selected VM. See "Starting a Virtual Machine" on page 156 .
 Boot From CD	Boots a VM from the selected virtual CD. See "Booting from a Virtual CD" on page 171 .
The following action is available if the everRun Enterprise software has removed the VM from service and powered it off because an excessive failure rate.	
 Reset Device	<p>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" on page 176.</p> <p>When a VM crashes, the everRun software automatically restarts it, unless it has fallen below its MTBF threshold. If the VM is below the MTBF threshold, the everRun software leaves it in the crashed state. If necessary, you can click Reset Device to restart the VM and reset the MTBF counter.</p>

Related Topics

["Managing the Operation of a Virtual Machine" on page 156](#)

["The Virtual Machines Page" on page 84](#)

["Using the everRun Availability Console" on page 51](#)

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		everRun Enterprise software is installing the boot volume for a new VM.
	 stopped	Start Config Boot From CD Remove	VM has been shut down or powered off.
	 booting	Console Power Off	VM is starting.
	 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 Shutdown 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 mes-

State	Activity	Enabled Actions	Description
			sages 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" on page 176.

Related Topics

["Managing the Operation of a Virtual Machine" on page 156](#)

["The Virtual Machines Page" on page 84](#)

["Using the everRun Availability Console" on page 51](#)

The Volumes Page

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

- **Name**
- **Size**
- **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**) 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 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 everRun System" on page 167](#)).
- Remove a volume by clicking the **Remove** icon. Note, though, that the **Remove** icon does not appear 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" on page 164](#)
- ["Creating a Volume in a Virtual Machine" on page 163](#)
- ["Detaching a Volume from a Virtual Machine" on page 165](#)
- ["Removing a Volume from a Virtual Machine" on page 166](#)

Related Topics

["The everRun Availability Console" on page 51](#)

["Using the everRun Availability Console" on page 51](#)

The Storage Group Page

The **Storage Group** page displays information about disks in the everRun Enterprise system. To open this page, click **Storage Group** in the left-hand navigation panel of the everRun Availability Console.

To view information about the storage group, click the storage group name in the top pane of the **Storage Group** page. The bottom pane displays additional information about the storage group.

You can use the **Storage Group** page to view information about the storage group, including name, size used, size, and number of volumes. You can also add a description of the storage group using the **Description** tab in the bottom pane.



Caution: The everRun software automatically synchronizes disks on the secondary physical machine (PM) with disks on the primary PM, when, for example, you change disks or when you upgrade or restore PMs. During synchronization of volumes between PMs, a busy icon (🔄) appears on **System** and **Volumes** in the left-hand navigation panel. Do not remove either PM during synchronization.

For more information about storage and everRun systems, see ["everRun Enterprise Storage Architecture" on page 15](#).

Related Topics

["The everRun Availability Console" on page 51](#)

["Using the everRun Availability Console" on page 51](#)

The Networks Page

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

To manage a specific network, click the name of a network 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 information about the network.

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

- ["Connecting Additional Networks" on page 48](#).
- 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.

For additional information on networks, see the following topic:

- ["General Network Requirements and Configurations" on page 24](#)
- ["SplitSite Network Requirements" on page 27](#)



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 everRun Availability Console" on page 51](#)

["Using the everRun Availability Console" on page 51](#)

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 your everRun Enterprise system. To open this page, click **Virtual CDs** in the left-hand navigation panel of the everRun Availability 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" on page 169](#)
- ["Removing a Virtual CD" on page 172](#)
- ["Renaming a Virtual CD" on page 172](#)
- Adding a description for each volume, in the **Description** text box

To complete other VCD management tasks, see ["Managing Virtual CDs" on page 169](#).

Related Topics

["Using the everRun Availability Console" on page 51](#)

The Upgrade Kits Page

The **everRun Upgrade Kits** page enables you to upload and manage software kits that you use to upgrade your system to newer versions of everRun Enterprise. To open the **Upgrade Kits** page, click **Upgrade Kits** in the left-hand navigation panel in the everRun Availability Console.

For information about upgrading your everRun software, see ["Upgrading everRun Enterprise Software" on page 95](#).

Related Topics

["The everRun Availability Console" on page 51](#)

["Using the everRun Availability Console" on page 51](#)

The Users & Groups Page

Use the **Users & Groups** page to add, modify, or remove user accounts on your everRun system. To open this page, click **Users & Groups** in the left-hand navigation panel of the everRun Availability 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" on page 92](#).

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

Related Topics

["Managing Local User Accounts" on page 92](#)

Managing Local User Accounts

You add, edit, or remove users, specify passwords, and assign ["User Roles" on page 93](#) to local-user accounts on the **User & Groups** page in the everRun Availability Console.

Local user accounts reside on the everRun 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.

To add a user account

1. In the lower left-hand panel, select **Users & Groups**.
2. In the top pane, click **Add**.
3. In the **Role** drop-down window, select **Administrator**, **Platform Manager**, or **Read-only**.
4. Provide values for the **User Name**, **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.
5. Click **Save**.

To edit a user account

1. In the lower left-hand panel, select **Users & Groups**.
2. In the top pane, click **Edit**.
3. To change a user's role, in the **Role** drop-down window, select **Administrator**, **Platform Manager**, or **Read-only**.
4. Click **Save**.

To remove a user account

1. In **Users & Groups**, select the account to remove.
2. Click **Remove**.
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.

User Roles

- **Administrator:** full system administrator privileges
- **Platform Manager:** system administrator privileges except for adding, removing, and modifying users
- **Read-only:** ability to view but not to change system configuration or to install system software

4

Chapter 4: Upgrading everRun Enterprise Software

This topic describes how to upgrade everRun Enterprise software.



Prerequisite: All PMs and VMs in must be in good health prior to performing an upgrade of everRun Enterprise software. Before starting an upgrade, examine the everRun Availability Console to verify that there are no alerts indicating PM or VM problems.

To upload an upgrade kit

1. In the everRun Availability 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 **everRun Enterprise - Kit Upload Wizard**.
3. In the **everRun Enterprise - 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** or **Finish** (they perform the same function). The message **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.
5. After the upload is complete, the message **Kit uploaded successfully**. Click **OK to close Wizard** appears. Click **OK** to close the wizard.

The **Upgrade Kits** page now lists the state and version number of the upgrade kit. The **Upgrade** and **Delete** buttons also appear with the **Add a Kit** button.

6. If more than one upgrade kit is loaded, select it.
7. Click **Upgrade** to upgrade the everRun system.

First, the everRun software upgrades and reboots the secondary PM. The newly upgraded PM becomes primary, and then the everRun software upgrades and reboots the other PM.



Note: This procedure also updates the AVCLI software on the everRun Enterprise system. If you have installed AVCLI on a remote management computer, you must manually upgrade AVCLI to the most recent version on the remote computer. You can obtain AVCLI software from the **Drivers and Tools** section of the **everRun Enterprise Downloads and Support** page at <http://www.stratus.com/go/support/everrun>. For information about how to manually install AVCLI on a remote computer see "[AVCLI Command Overview](#)" on page 194.

5

Chapter 5: Migrating From Non-everRun Enterprise Systems

If you are migrating from an everRun MX system or an Avance unit to an everRun Enterprise system, and you want to transfer virtual machines (VMs) from the other system, see ["Creating and Migrating Virtual Machines" on page 125](#).

To learn more about migrating your system-wide configuration to an everRun Enterprise system, start with one of the following topics, depending on your needs:

- ["Planning to Migrate from an everRun MX System" on page 97](#)

Use this planning information to consider the system-wide configuration and settings that are affected when you migrate an everRun MX system and its VMs to an everRun Enterprise system.

- ["Planning to Migrate from an Avance Unit" on page 100](#)

Use this planning information to consider the system-wide configuration and settings that are affected when you migrate an Avance unit and its VMs to an everRun Enterprise system.

Related Topics

["Planning" on page 21](#)

["Software Installation" on page 32](#)

["Post-Installation Tasks" on page 45](#)

["Creating and Migrating Virtual Machines" on page 125](#)

Planning to Migrate from an everRun MX System

If you have an existing everRun MX system, this topic describes some items to consider when migrating to an everRun Enterprise system.

For all systems, see ["Creating and Migrating Virtual Machines" on page 125](#) for information about migrating your virtual machines (VMs) to the everRun Enterprise system.



Note: For best results, contact your authorized Stratus service representative for assistance in evaluating and performing upgrades from any non-everRun Enterprise system.

Platform Requirements

Whether you reuse the existing everRun MX hardware or migrate to new hardware, the platform must meet the minimum system requirements for everRun Enterprise systems, as described in ["Physical Machine System Requirements" on page 317](#).

Multi-node XenServer pools are supported by everRun MX, but only two-node configurations are supported for everRun Enterprise systems.

Planned Outage

The considerations in this help topic assume that an outage can be tolerated throughout the migration process. If you have minimum downtime requirements, contact your authorized Stratus service representative for assistance.

Guest Operating System Support

Verify that the Windows guest operating system running in each of the everRun MX virtual machines is supported by the everRun Enterprise software. See ["Compatible Guest Operating Systems" on page 316](#).

Also, verify that each Windows guest operating system is supported by the migration process (as described in ["Migrating a Physical Machine or Virtual Machine to the everRun Enterprise System" on page 128](#)) or the import process (as described in ["Importing an OVF File from an everRun MX System to the everRun Enterprise System" on page 135](#)).

Network Preparation

Prepare the platform network and the networking environment to conform to the everRun Enterprise requirements. See ["General Network Requirements and Configurations" on page 24](#).

Management Network Access

The XenServer management network becomes the everRun Enterprise business network. As in everRun MX, the management console (everRun Availability Console) is accessed over this network.

Bonded network interfaces are recommended for the XenServer management network, but they are not supported for the everRun Enterprise management network.

In everRun MX, each node in the XenServer pool has an IPv4 address associated with it. The same is true for an everRun Enterprise system, but a **system IP** address is also required, and it must be a static address (not DHCP). This system IP address provides access to the everRun Availability Console; it is failed-over between everRun Enterprise nodes as necessary by the everRun Enterprise software.

Availability Link Networks

The A-Link networks that were in use in everRun MX will continue to be the A-Link networks on the everRun Enterprise system. In everRun MX, the A-Links could have networking interfaces in each node that were not on the same subnet, but this is not possible on an everRun Enterprise system. For each of the two possible A-Links, the network interfaces associated with them in each node must exist in the same local network, because IPv6 link local addresses are used to identify them.

Two 10-Gb networks are recommended for the A-Links.

It is not required that the A-Link connections be point-to-point (that is, they can be on a switched network).

Private Network

The everRun Enterprise private network must be identified. Only one everRun Enterprise system can be installed and operational on the private network at any one time, so it is recommended that the private network be a point-to-point connection between the two everRun Enterprise nodes.

On the everRun Enterprise system, it is typical to share one of the A-Links for the private network when at least one of the A-Link networks is connected point-to-point.

A 10-Gb network is recommended for the private network.

Business Networks

All networks that are not the private network or an A-Link network can also be business networks (that is, networks available for use by the VMs). The management network can be used simultaneously as a business network.

Storage Considerations

everRun MX supported both external storage and redundant-path storage. Neither of these configurations is supported on an everRun Enterprise system.

everRun MX allowed the storage to be configured into multiple volume groups. The everRun Enterprise software automatically creates a single storage group out of all the available storage.

For physical storage requirements, see ["Storage Requirements" on page 23](#).

Quorum Support

Prior to everRun MX 6.2, the quorum servers could be available only over the A-Links. As of everRun MX 6.2, the quorum servers could be available over any network in the XenServer pool. On everRun Enterprise systems, the quorum servers must be available over the business network, which is configured with an IPv4 address and required for quorum.

The Preferred Quorum Server should be configured as the first Quorum Server and the Alternate Quorum Server should be configured as the second Quorum Server in the everRun Availability Console.

Installing everRun Enterprise

After the nodes in the everRun Enterprise system are configured, you can install and configure the everRun Enterprise software, as described in ["Software Installation" on page 32](#).

Migrating Virtual Machines

Using either the P2V client migration process or the OVF import process, migrate the VMs to the everRun Enterprise system. For an overview of each process, see ["Creating and Migrating Virtual Machines" on page 125](#).

Planning to Migrate from an Avance Unit

If you have an existing Avance unit, this topic describes some items to consider when migrating to an everRun Enterprise system.

For all systems, see ["Creating and Migrating Virtual Machines" on page 125](#) for information about migrating your virtual machines (VMs) to the everRun Enterprise system.



Note: For best results, contact your authorized Stratus service representative for assistance in evaluating and performing upgrades from any non-everRun Enterprise system.

Platform Requirements

Whether you reuse the existing Avance hardware or migrate to new hardware, the platform must meet the minimum system requirements for everRun Enterprise systems, as described in ["Physical Machine System Requirements" on page 317](#).

Planned Outage

The considerations in this help topic assume that an outage can be tolerated throughout the migration process. If you have minimum downtime requirements, contact your authorized Stratus service representative for assistance.

Guest Operating System Support

Verify that the Windows or Linux guest operating system running in each of the Avance VMs is supported by the everRun Enterprise software. See ["Compatible Guest Operating Systems" on page 316](#).

Also, verify that each guest operating system is supported by the migration process (as described in ["Migrating a Physical Machine or Virtual Machine to the everRun Enterprise System" on page 128](#)) or the import process (as described in ["Importing an OVF File from an everRun MX System to the everRun Enterprise System" on page 135](#)).

Network Preparation

Prepare the platform network and the networking environment to conform to the everRun Enterprise system requirements. See ["General Network Requirements and Configurations" on page 24](#).

Management Network Access

The same network that was used to access the Avance Management Console will be used for the everRun Availability Console.

In Avance, the nodes were available on the management network for management only through the IPv4 system address, which could be failed-over to either node in the system. The everRun Enterprise software uses the same system address, but it also requires separate IPv4 addresses for each node in the same subnet as the system IP address.

Availability Link Networks

Avance had no Availability Links; therefore, these networks must be added to the hardware configuration.

Two 10-Gb networks are recommended for the A-Links.

It is not required that the A-Link connections be point-to-point (that is, they can be on a switched network).

Private Network

The same network used as the private network on the Avance unit can be used for the private network on the everRun Enterprise system.

Only one everRun Enterprise system can be installed and operational on the private network at any one time, so it is recommended that the private network be a point-to-point connection between the two everRun Enterprise nodes.

It is typical to share one of the A-Links for the private network when at least one of the A-Link networks is connected point-to-point.

A 10-Gb network is recommended for the private network.

Business Networks

All networks that are not the private network or an A-Link network can also be business networks (that is, networks available for use by the VMs). The management network can be used simultaneously with the business network.

Storage Considerations

Avance unit storage can be used as-is in the everRun Enterprise system; however, there can be only one storage group. For physical storage requirements, see ["Storage Requirements" on page 23](#).

Installing everRun Enterprise

After the nodes in the everRun Enterprise system are configured, you can install and configure the everRun Enterprise software, as described in ["Software Installation" on page 32](#).

Migrating Virtual Machines

Using either the P2V client migration process or the OVF import process, migrate the VMs to the everRun Enterprise system. For an overview of each process, see ["Creating and Migrating Virtual Machines" on page 125](#).

6

Chapter 6: Managing Logical Disks

Manage logical disks using the everRun Availability Console. For overview information see ["Logical Disk Management" on page 103](#) as well as ["Logical Disks and Physical Disks" on page 15](#).

To perform tasks, see the following:

- ["Responding to a Failed Logical Disk" on page 104](#)
- ["Activating a New Logical Disk" on page 106](#)

Logical Disk Management

In an everRun Enterprise system, use the everRun Availability Console to manage logical disks by activating a new logical disk and responding to a failed logical disk.

In some cases, you need to activate a new logical disk even though the everRun Enterprise software automatically recognizes new logical disks that the RAID controller presents to the operating system. For information, see ["Activating a New Logical Disk" on page 106](#).

You need to respond to alerts regarding missing or failed logical disks. everRun Enterprise software may detect a logical disk failure when a physical disk is removed or fails. The everRun Enterprise software then generates an alert that appears on the DASHBOARD. The following alerts are examples:

- **System has missing or failed Logical Disks.**
- **Logical Disk - 1 on PM node1 has failed.**

On the **Physical Machines** page of the everRun Availability Console, the Storage tab for each PM identifies logical disks that have failed. For information, see ["The Physical Machines Page" on page 80](#).

When a logical disk has failed, system storage is frozen. You cannot allocate new volumes until you have responded to the alert. Your response might require using the RAID controller BIOS or the **Repair** button on the masthead. For information, see ["Responding to a Failed Logical Disk" on page 104](#)

Related Topics

["Logical Disks and Physical Disks" on page 15](#)

["The everRun Availability Console" on page 51](#)

Responding to a Failed Logical Disk

When everRun Enterprise software detects a missing or broken logical disk, it displays a failed logical disk alert on the DASHBOARD page of the everRun Availability Console. (For examples of alerts, see ["Logical Disk Management" on page 103](#).) You can also view the alert on the ALERT HISTORY page. The everRun Availability Console continues to display the alert until you respond to the problem using one of the following methods, as appropriate for your situation:

- If a physical disk has been pulled, reinsert the appropriate physical disk. In this case, the physical machine restores the disk and you may need to use RAID controller software to complete the logical disk restoration.
- If a logical disk is broken or missing, you can attempt to use RAID controller software to recover it. If you are able to use RAID controller software to restore the logical disk to service, the everRun software will detect the restored logical disk and start using its data
- If a logical disk is broken or missing, and you cannot recover the logical disk using RAID controller software (for example, a failed physical disk needs to be replaced), click the **Repair** button in the masthead to complete the repair. After clicking the **Repair** button, the everRun software:
 - Dismisses the alert.
 - Evacuates all failed logical disks.
 - Removes all failed logical disks from their storage groups.
 - Attempts to repair any volumes that had been using the failed logical disks.

Caution:

1. Clicking the **Repair** button removes all data on failed logical disks.
2. Repairing storage causes virtual machines (VMs) that are using failed logical disks to become simplex until repair is complete.
3. In some configurations, if you need to repair a logical disk that is the boot disk, you may need to reconfigure the RAID controller to boot from one of the remaining logical disks. Any logical disk that is not affected by the failed disk is able to boot the server. The everRun Enterprise software mirrors the boot files for each node, in order to maximize overall availability. However, some systems may be able to boot from only the predefined boot logical disk in the RAID controller, and may be unable to boot from an alternate logical disk, if the predefined boot logical disk is present but not bootable. After the node has recovered and the logical disk with the replacement drive has been brought up to date, you should restore the boot device to the original value in the RAID controller.

**To repair a failed logical disk**

1. Click the **Repair** button that appears in the masthead of the everRun Availability Console.
2. Click **Yes** in the Confirm message box if you want to continue with the repair.

After you click the **Repair** button, the everRun software attempts to repair all broken volumes by migrating data to other logical disks. When other logical disks have enough space for the data, the everRun software can successfully complete the repair. When other logical disks do not have enough space for the data, the everRun software generates the alert **Not enough space for repair**. In this case, you need to add more storage to the storage group by creating new logical disks or by deleting some existing volumes.

When enough space for the data exists, the everRun software automatically re-mirrors broken volumes.

After the repair is complete, use RAID controller software to remove the failed logical disk and to create a new logical disk. everRun software automatically recognizes the new logical disk and brings it into service if the disk does not contain data. If the disk contains data, the DASHBOARD displays the message **Logical Disk - *n* on PM node *n* is foreign and should be activated or removed**. To activate the logical disk, see ["Activating a New Logical Disk" on page 106](#).

Related Topics

["Logical Disks and Physical Disks" on page 15](#)

["The everRun Availability Console" on page 51](#)

Activating a New Logical Disk

In an everRun Enterprise system, the RAID controller creates logical disks from the system's physical disks. The everRun software is able to access logical disks that the RAID controller presents to the operating system. When the everRun software recognizes a new logical disk, it performs one of the following actions:

- If the logical disk does not contain data, the everRun software brings the logical disk into service.
- If it is a known logical disk that was not evacuated, the everRun software starts using the logical disk and its data.
- If the disk contains unknown data, the DASHBOARD displays the message **Logical Disk - *n* on PM node*n*** is foreign and should be activated or removed. In this case, you can activate or remove the disk now, or you can do nothing now, but later activate or remove the disk.



Caution: Activating the logical disk causes all data on the disk to be lost.

To activate a new logical disk

1. Click **Physical Machines** in the left-hand navigation panel.
2. On the **Physical Machines** page, select either **node0** or **node1** in the top pane.
3. On the **Physical Machines** page, click the **Storage** tab in the bottom pane.
4. In the **Action** column, click the **Activate Foreign** button in the Action column to activate the corresponding logical disk.
5. When the **Confirm** message box appears, click **Yes** to confirm activating the logical disk. Activating the logical disk causes all data on the disk to be lost.

The everRun software partitions the new logical disk, adds it to the Initial Storage Group, and begins to use the disk.

Related Topics

["Responding to a Failed Logical Disk" on page 104](#)

["Logical Disk Management" on page 103](#)

["Logical Disks and Physical Disks" on page 15](#)

["The everRun Availability Console" on page 51](#)

7

Chapter 7: Managing Physical Machines

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

You view and manage PMs using the **Physical Machines** page of everRun Availability Console; for information, see ["The Physical Machines Page" on page 80](#).

Many of the tasks you perform from the **Physical Machines** page require maintenance mode; for information, see ["Maintenance Mode" on page 109](#).

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

- ["Shutting Down a Physical Machine" on page 112](#)
- ["Rebooting a Physical Machine" on page 111](#)
- ["Load Balancing" on page 113](#)

To troubleshoot a PM by recovering a failed PM or resetting MTBF for a failed machine, see ["Troubleshooting Physical Machines" on page 113](#).

To perform maintenance tasks, see ["Maintaining Physical Machines" on page 179](#).

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).

When one PM enters maintenance mode, the PM migrates the virtual machines (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.

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

Cautions:



1. An everRun system is not fault tolerant 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. Avoid entering both PMs into maintenance mode at the same time. To keep the virtual machines running, at least one PM must be up and running normally. (If you need to shut down the entire everRun system, see ["Shutting Down a Physical Machine" on page 112.](#))



Note: 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.

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 physical machine from the **Physical Machines** page.
2. Click **Finalize**, which exits the PM from maintenance mode.

Related Topics

["The everRun Availability Console" on page 51](#)

["Managing Physical Machines" on page 109](#)

["Physical Machines and Virtual Machines" on page 8](#)

["The Physical Machines Page" on page 80](#)

["The Virtual Machines Page" on page 84](#)

Physical Machine Management Actions

You can perform the following physical machine management actions:

- ["Rebooting a Physical Machine" on page 111](#)
- ["Shutting Down a Physical Machine" on page 112](#)
- ["Load Balancing" on page 113](#)

Rebooting a Physical Machine

Reboot a physical machine (PM) to restart its everRun Enterprise software, and optionally exit the PM from maintenance mode. (If you need to reboot both PMs in the everRun system, see ["Rebooting the System" on page 55.](#))

To reboot a PM

1. Determine which PM (node0 or node1) you want to reboot.
2. In the everRun Availability Console, click **Physical Machines** in the left-hand navigation panel.
3. 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)**.
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**.

Related Topics

["Maintenance Mode" on page 109](#)

["The everRun Availability Console" on page 51](#)

["Managing Physical Machines" on page 109](#)

["The Physical Machines Page" on page 80](#)

Shutting Down a Physical Machine

Shut down a physical machine (PM) to stop it from running when you need to service or replace the PM.

Use this procedure to shut down one and only one PM.

Cautions:



1. Using this procedure to shut down both PMs will result in data loss. If you need to stop both PMs, shut down the everRun system (which also shuts down the virtual machines (VMs)), as described in ["Shutting Down the System" on page 56](#).
2. The everRun Enterprise system is not fault tolerant when you shut down a PM. For continuous uptime, bring an offline PM back into service as soon as possible.

To shut down a PM, you must place the PM into maintenance mode, which migrates any virtual machines running on that PM to the remaining PM.

To shut down a PM

1. Determine which PM you want to shut down.
2. In the everRun Availability Console, click **Physical Machines** in the left-hand navigation panel.
3. 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)**.
4. After the PM displays **running (in Maintenance)**, click **Shutdown**.



Caution: If the PM does not turn off after you click **Shutdown**, you must manually power off the PM, though doing so destroys its memory state. **Use this action only as a last resort.**

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

Related Topics

["Maintenance Mode" on page 109](#)

["The everRun Availability Console" on page 51](#)

["Managing Physical Machines" on page 109](#)

["The Physical Machines Page" on page 80](#)

Load Balancing

HA Load Balancing distributes VMs across both PMs to improve performance and availability. Load balancing is configured per VM and is enabled automatically on everRun Enterprise systems.

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, the alert is generated. The alert appears on the Dashboard, and a Load Balancing notification appears on the masthead.

Click **Load Balance** 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 nodeN.** 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 under the **Current PM** tab 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 VMs are currently at PM 0 and that PM 1 is the preference.



everRun 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" on page 174](#)

Troubleshooting Physical Machines

The following topics describe troubleshooting procedures for PMs:

- ["Recovering a Failed Physical Machine" on page 114](#)
- ["Resetting MTBF for a Failed Physical Machine" on page 116](#)

Recovering a Failed Physical Machine

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

To recover a PM, you must reinstall the everRun Enterprise release that the PM has been running using an install ISO. Recovering a failed PM, though, is different from installing the software for the first time. The recovery preserves all data, but it recreates the /boot and root file systems, re-installs CentOS and the everRun Enterprise software, and attempts to connect to an existing system.



Note: If you need to repair or replace the PM, see ["Replacing Physical Machines, Motherboards, NICs, or RAID Controllers" on page 182](#), which requires the **Replace** operation of a node in maintenance mode.

Prerequisites:

1. Determine which PM you need to recover.
2. Obtain installation software for the everRun Enterprise release that the PM has been running by using one of the following methods:
 - Download an install ISO from your authorized Stratus service representative.
 - Extract an install ISO into the current working directory from the most recently used upgrade kit by executing a command similar to the following (*x.x.x.x* is the release number and *nnn* is the build number):

```
tar -xzvf everRun_Enterprise_upgrade-x.x.x.x-nnn.kit *.iso
```

After you obtain the correct install ISO, save it or burn it to a DVD. See ["Obtaining everRun Enterprise Software" on page 34](#).

3. Check that a monitor and keyboard are connected to the PM that you are recovering.
4. Check that Ethernet cables are connected from the PM you are recovering to the network or directly to the other PM, if the two everRun system PMs are in close proximity. The Ethernet cable should connect from the first embedded port on the PM you are recovering or from an option (that is, add-on or expansion) port if the PM does not have an embedded port.

To recover a PM

1. Manually power on the PM that you want to recover. As the PM powers on, enter the BIOS and set the Optical Drive as the first boot device.
2. Either mount the ISO image or insert the DVD into the PM.
3. At the **Welcome** screen, select **Recover PM, Join system: Preserving data** and press **Enter**.
4. When prompted, respond to the **Select interface for private Physical Machine connection**, and then respond to the prompt **Select interface for managing the system (ibiz0)**.
5. When prompted to configure **ibiz0**, select **Automatic configuration via DHCP** or **Manual Configuration (Static Address)**. (The installation software configures **priv0** automatically.)
6. When installation is complete, the PM ejects the install DVD (if used) and reboots.

7. As the PM boots, you can view its activity on the **Physical Machines** page of the everRun Availability Console. The **Activity** column displays the PM as **recovery (in Maintenance)** and then **running** after recovery is complete.

Related Topics

["Maintenance Mode" on page 109](#)

["Managing Physical Machines" on page 109](#)

["The everRun Availability Console" on page 51](#)

["The Physical Machines Page" on page 80](#)

Resetting MTBF for a Failed Physical Machine

Reset the mean time between failures (MTBF) counter for a physical machine (PM) to attempt to restart a failed PM.

If a PM crashes, the everRun Enterprise software automatically restarts it, unless it has fallen below its MTBF threshold. If the PM is below the MTBF threshold, the everRun software leaves the machine powered off. If necessary, you can reset the MTBF counter and restart the PM.



Caution: Do not reset the MTBF counter unless instructed to do so by your authorized Stratus service representative, as doing so may affect the fault tolerance of your system.



Note: The **Reset Device** button is displayed only if the PM falls below its MBTF threshold.

To reset the MTBF counter for a PM

1. Determine which PM has the MTBF counter you want to reset.
2. In the everRun Availability Console, click **Physical Machines** in the left-hand navigation panel.
3. 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)**.
4. After the PM displays **running (in Maintenance)**, click **Reset Device**.

Related Topics

["Maintenance Mode" on page 109](#)

["Managing Physical Machines" on page 109](#)

["The everRun Availability Console" on page 51](#)

["The Physical Machines Page" on page 80](#)

8

Chapter 8: 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 everRun Availability Console, which you access as described in ["The Virtual Machines Page" on page 84](#). To perform specific management tasks, see the following topics.

To manage the operational state of a VM, see:

- ["Starting a Virtual Machine" on page 156](#)
- ["Shutting Down a Virtual Machine" on page 156](#)
- ["Powering Off a Virtual Machine" on page 157](#)
- ["Opening a Virtual Machine Console Session" on page 158](#)

To create or configure a VM, see:

- ["Planning Virtual Machine Resources" on page 120](#) (virtual CPUs, memory, storage, and networks)
- ["Creating and Migrating Virtual Machines" on page 125](#)
- ["Managing Virtual CDs" on page 169](#)
- ["Configuring Windows-based Virtual Machines" on page 152](#)
- ["Configuring Linux-based Virtual Machines" on page 154](#)
- ["Managing Virtual Machine Resources" on page 160](#)

To complete advanced tasks, see:

- ["Assigning a Specific MAC Address to a Virtual Machine" on page 173](#)
- ["Selecting a Preferred PM for a Virtual Machine" on page 174](#)
- ["Changing the Protection Level for a Virtual Machine \(HA or FT\)" on page 174](#)
- ["Configuring the Boot Sequence for Virtual Machines" on page 175](#)
- ["Resetting MTBF for a Failed Virtual Machine" on page 176](#)
- ["Locating a Dump File in a Virtual Machine" on page 177](#)

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" on page 120](#)
- ["Planning Virtual Machine Memory" on page 122](#)
- ["Planning Virtual Machine Storage" on page 123](#)
- ["Planning Virtual Machine Networks" on page 124](#)

Planning Virtual Machine vCPUs

Allocate virtual CPUs (vCPUs) to assign computing resources to a virtual machine (VM) on your everRun Enterprise 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 an everRun Enterprise 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 another PM 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 available on the everRun Enterprise system minus the number of vCPUs allocated to the everRun system software (which you can set to 2 or 4 vCPUs, as described in ["Configuring System Resources" on page 69](#)). For example, if the total number of vCPUs is 8, and you allocate 2 vCPUs to the system software, there are 6 vCPUs that can be assigned to running VMs without over-provisioning the system.

- The maximum number of vCPUs you can allocate to any one VM is the total number of vCPUs in the system. Each VM consumes its configured number of vCPUs plus an additional 2 vCPUs for overhead.
- 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" on page 161](#)), 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 everRun Availability Console (see ["The System Page" on page 55](#)) indicates the total number of vCPUs, the number of vCPUs allocated to the everRun system software, the number of vCPUs consumed by running VMs, and the number of free vCPUs.
- The everRun 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 everRun 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" on page 22](#)

["Creating and Migrating Virtual Machines" on page 125](#)

["Managing Virtual Machine Resources" on page 160](#)

Planning Virtual Machine Memory

Allocate memory to assign physical memory to a virtual machine (VM) on your everRun Enterprise 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 everRun system (see ["Memory Requirements" on page 24](#)) minus the memory allocated to the everRun system software (which you can set to 1, 2, or 4 gigabytes (GB), as described in ["Configuring System Resources" on page 69](#)). For example, if the total amount of memory is 16 GB, and you allocate 2 GB to the system software, there are 14 GB of memory available to the VMs.
- You can provision a single VM with memory up 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 everRun Availability Console (see ["The System Page" on page 55](#)) indicates the total amount of memory, the memory allocated to the everRun system software, the memory consumed by running VMs, and the amount of free memory. Use this page to verify your memory allocations.
- The everRun 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

["Memory Requirements" on page 24](#)

["Creating and Migrating Virtual Machines" on page 125](#)

["Managing Virtual Machine Resources" on page 160](#)

Planning Virtual Machine Storage

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

When you install the everRun software, it forms a storage group from the available space on all logical disks. You allocate VM volumes and virtual CDs (VCDs) from this storage group, the assignment of which can have a dramatic impact on system performance and your ability to fully utilize available storage capacity.

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

- Observe storage maximums

The everRun 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 everRun system. The system prevents you from creating a volume for a VM from a storage group that does not have enough space.

- Minimize stranded storage

Ensure that each PM has the same amount of storage capacity. If one PM has more storage than the other, only the minimum amount of space is available in the storage group. For example, if one PM has 3 terabytes (TB) of storage and another PM has 2 TB of storage, the total amount of storage is 2 TB (least storage of either PM).

- Leave space for additional VCDs

Leave at least 5 GB of free space in a storage group 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.

Related Topic

["Storage Requirements" on page 23](#)

["Creating and Migrating Virtual Machines" on page 125](#)

["Managing Virtual Machine Resources" on page 160](#)

Planning Virtual Machine Networks

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

When you install the everRun software, it binds pairs of physical network ports across the two physical machines (PMs) to form redundant virtual networks. When you create or reprovision VMs on your everRun system, you connect the VMs to these virtual networks instead of the physical network ports.

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 everRun 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" on page 24](#)

["Creating and Migrating Virtual Machines" on page 125](#)

["Managing Virtual Machine Resources" on page 160](#)

Creating and Migrating Virtual Machines

Generate a new virtual machine (VM) on the everRun Enterprise 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 from an existing everRun MX or Avance VM.

To create a new VM (without an existing source VM or PM), see ["Creating a New Virtual Machine" on page 125](#).

To migrate or import a system from a non-everRun Enterprise source, see ["Migrating From Non-everRun Enterprise Systems" on page 97](#) for considerations, and then see one of the following topics depending on your needs:

- ["Migrating a Physical Machine or Virtual Machine to the everRun Enterprise System" on page 128](#) (most VMs or PMs, including everRun MX- and Avance-based VMs)

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

- ["Importing an OVF File from an everRun MX System to the everRun Enterprise System" on page 135](#) (everRun MX-based VMs only)

Use XenConvert to export a VM from the everRun MX system to OVF and Virtual Hard Disk (VHD) files on a network share, and then use everRun Availability Console to import those files to the everRun Enterprise system.

- ["Importing an OVF File from an Avance System to the everRun Enterprise System" on page 143](#) (Avance-based VMs only)

Use Avance Management Console to export a VM from the Avance Unit to OVF and raw tar hard disk files on a management PC or network share, and then use everRun Availability Console to import those files to the everRun Enterprise system.

Creating a New Virtual Machine

Create a new virtual machine (VM) to install a guest operating system on your everRun Enterprise system. (You can also migrate an existing VM or physical machine (PM), as summarized in ["Creating and Migrating Virtual Machines" on page 125.](#))

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" on page 120.](#)
- Create a bootable virtual CD (VCD) of the Windows or Linux installation media, as described in ["Creating a Virtual CD" on page 169.](#) The bootable VCD must be a single CD or DVD. Multiple CDs or DVDs are not supported.
- Ensure that both PMs of the everRun Enterprise system are online; otherwise, the system cannot properly create the VM.

To create a new VM

1. On the **Physical Machines** page (see ["The Physical Machines Page" on page 80](#)), 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" on page 84](#)), 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 everRun Availability Console.
 - b. Select the level of protection to use for the VM:
 - **High Availability (HA)** – Provides basic failover and recovery, with some faults requiring an (automatic) VM reboot for recovery. Use HA for applications that can tolerate some downtime and that do not need the downtime protection that FT provides.

- **Fault Tolerant (FT)** – Transparently protects an application by creating a redundant environment for a VM running across two physical machines. Use FT for applications that need greater downtime protection than HA provides.

For more information about these levels of protection, see ["Modes of Operation" on page 11](#).

- c. Select the **Boot VCD** that contains the operating system you want to install.
 - d. Click **Next**.
4. On the **Volumes** page:
- a. Type the **Name** of the boot volume as it will appear in the everRun Availability Console.
 - b. Type the **Size** of the volume to create in gigabytes (GB). (For information about allocating storage, see ["Planning Virtual Machine Storage" on page 123](#).)
 - c. Select the **Storage Group** in which to create the volume.
 - d. If applicable, create additional data volumes by clicking **Add New Volume**. Specify the **Name**, **Size**, and **Storage Group** 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" on page 163](#).)
 - e. To continue, click **Next**.
5. On the **Networks** page, select the shared networks to assign to the VM. For more information, see ["Planning Virtual Machine Networks" on page 124](#). To continue, click **Next**.
6. On the **vCPUs and Memory** page, specify the number of **vCPUs** and the amount of **Memory** to assign to the VM. For more information, see ["Planning Virtual Machine vCPUs" on page 120](#) and ["Planning Virtual Machine Memory" on page 122](#). 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**.
8. If applicable, observe the progress of the installation of the operating system 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" on page 152](#)
- ["Configuring Linux-based Virtual Machines" on page 154](#)



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

["Renaming a Virtual Machine" on page 159](#)

["Removing a Virtual Machine" on page 160](#)

["Creating and Migrating Virtual Machines" on page 125](#)

["Managing Virtual Machine Resources" on page 160](#)

["Managing the Operation of a Virtual Machine" on page 156](#)

Migrating a Physical Machine or Virtual Machine to the everRun Enterprise System

Migrate a physical machine (PM) or virtual machine (VM) to transfer it over the network to a new VM on the everRun Enterprise system. (You can also import an Open Virtualization Format (OVF) file to the everRun Enterprise system as described in ["Importing an OVF File from an everRun MX System to the everRun Enterprise System" on page 135](#) or ["Importing an OVF File from an Avance System to the everRun Enterprise System" on page 143](#).)

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



Caution: Consider backing up your source PM or VM before preparing to migrate it.

Notes:

- The migration process supports only PMs or VMs running CentOS/RHEL 6, Windows 7, Windows Server 2008, Windows Small Business Server 2011, or Ubuntu 12.04 or later.
- 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 everRun Enterprise system, your 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 process, as described below in **Troubleshooting**.
- For Ubuntu-based PMs or VMs, you must edit the `/boot/grub/grub.cfg` file and change the `gfxmode` parameter to `text` (for example, `set gfxmode=text`) in the source PM or VM before the migration; otherwise, the new VM's console hangs on the everRun Enterprise system. You can restore the original setting in the source PM or VM after the migration.
- Your source PM or VM must be shut down for the duration of the migration process. Consider scheduling a planned maintenance period for the migration.
- The time required for the PM or VM 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 everRun Enterprise system. For example, transferring a source system with one 20 GB boot disk over a 1 Gb/s network may take about 30 minutes.
- You can migrate multiple PMs or VMs at one time, but sharing network bandwidth may increase migration times.
- If you will continue to use the source PM or VM after the migration, remember to set a different MAC address and IP address for the new VM on the everRun Enterprise system.
- If the everRun Enterprise 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.





Prerequisite: Both PMs of the everRun Enterprise system must be online for the migration process to function properly. On the **Physical Machines** page of the everRun Availability 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 prepare for migrating a PM to the everRun Enterprise system

1. Download the P2V client ISO file from the **Drivers and Tools** section of the **everRun Enterprise Downloads and Support** page at <http://www.stratus.com/go/support/everrun>.
2. Burn the P2V client ISO file to a CD-ROM that you will use to boot the physical machine.
3. Insert the P2V client CD into the CD/DVD drive of the source PM.
4. Shut down the PM in preparation to boot the P2V client.

To prepare for migrating a VM to the everRun Enterprise system

1. Download the P2V client ISO file from the **Drivers and Tools** section of the **everRun Enterprise Downloads and Support** page at <http://www.stratus.com/go/support/everrun>.
2. 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.
3. Shut down the VM in preparation to boot the P2V client.

To migrate a PM or VM to the everRun Enterprise 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. Configure the network settings to use for the migration process:
 - a. Select an active network device, if more than one is present.
 - b. To specify static network settings, optionally clear the **Automatic configuration** check box and enter your **IP Address**, **Gateway**, and **DNS Servers** settings.
Otherwise, keep the default settings to use DHCP.
 - c. Click **Use these network settings**.
3. Specify the connection settings for the conversion server (the everRun Enterprise 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 everRun host operating system, as described in ["Accessing the Host Operating System" on page 18.](#))

4. Click **Connect**. The **Target properties** page is displayed.
5. Select **EverRun-FT** as the **Destination Profile**.
6. Enter the **Name** for the target VM that will be displayed in the everRun Availability Console. (The name must be different from any existing VMs on the everRun Enterprise system.)
7. The **Number of CPUs** and **Memory(MB)** values are automatically detected and completed, but optionally modify them if you want the VM on the everRun Enterprise system to have more CPUs or memory than the source PM or VM.
8. Select which **Fixed Storage** devices to include in the migration by activating the check box next to each device.

You must select at least one **Fixed Storage** device, including the boot volume. (Because the P2V client is a Linux-based utility, all devices are listed by Linux device names, where **sda** represents the boot volume.)

The P2V client automatically manages **Removable Media** and **Network Interfaces** for the migration. Only one CD/DVD drive and one network interface is transferred to the new VM on the everRun Enterprise system, regardless of your selections. You cannot change the CD/DVD configuration in the new VM, but you can provision additional network interfaces to the VM as needed after the migration.

9. When you are ready to migrate the PM or VM to the everRun Enterprise system, click **Convert**. (If you need to cancel the migration for any reason, see **Troubleshooting** below.)
10. 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 everRun Enterprise 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 everRun Enterprise system

1. Open the **Virtual Machines** page (see ["The Virtual Machines Page" on page 84](#)) in the everRun Availability 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](#)" on page 161. 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](#)" on page 158.)
5. For Windows-based VMs only, install the required VirtIO drivers (already installed on Linux-based systems):



Note: You must update two or more drivers, each of which requires a system reboot. When prompted, you can wait until the last driver is updated to restart the guest operating system.

- a. Open **Device Manager** in the guest operating system.
- b. Expand **Other devices** and right-click the **Ethernet Controller** device. Select **Update Driver Software** and follow the wizard to automatically search for and install the **Red Hat VirtIO Ethernet Adapter** driver. Repeat the driver update for each additional **Ethernet Controller** device.
- c. Expand **Storage controllers** and right-click a **SCSI controller** entry. Select **Update Driver Software** and follow the wizard to automatically search for and install the **Red Hat VirtIO SCSI controller** driver. Repeat the driver update for each additional **SCSI controller** device.
- d. Restart the guest operating system to load the updated drivers.



Note: On the **Virtual Machines** page and **Volumes** page of the everRun Availability Console, the **State** column may not accurately reflect the state of the VM or volume until the VirtIO drivers are properly installed.

6. Disable any guest operating system services that are unnecessary for operation on the everRun Enterprise 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

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

7. If necessary, update the network configuration settings in the guest operating system and restart it to enable the settings.

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

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 everRun Availability 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 everRun Enterprise 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" on page 71](#). The diagnostic file includes any server-side debugging messages from the migration process.

To recover when the new VM's console hangs on the everRun Enterprise system

For Ubuntu-based VMs, the VM console hangs in everRun Availability Console if you do not properly set the `gfxmode` parameter before the migration process (as described in **Notes**). If the VM console hangs, keep restarting the VM until the console opens properly in everRun Availability Console and then modify the `gfxmode` parameter to prevent subsequent issues.

For additional VM console troubleshooting, see ["Opening a Virtual Machine Console Session" on page 158](#).

To recover missing data volumes in the VM on the everRun Enterprise system

If your data volumes do not appear in the VM on the everRun Enterprise 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, from Avance (`/dev/xvda` through `/dev/xvdh`) to everRun Enterprise (`/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 everRun Enterprise system

If your network devices do not appear in the VM on the everRun Enterprise 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.

Related Topics

["Migrating From Non-everRun Enterprise Systems" on page 97](#)

["Creating and Migrating Virtual Machines" on page 125](#)

["Configuring Windows-based Virtual Machines" on page 152](#)

["Configuring Linux-based Virtual Machines" on page 154](#)

["Managing Virtual Machine Resources" on page 160](#)

["Managing the Operation of a Virtual Machine" on page 156](#)

Importing an OVF File from an everRun MX System to the everRun Enterprise System

Import an Open Virtualization Format (OVF) file from an everRun MX system if you want to transfer a VM to the everRun Enterprise system for deployment. (To migrate a physical machine (PM) or virtual machine (VM) to the everRun Enterprise system without using an OVF file, see ["Migrating a Physical Machine or Virtual Machine to the everRun Enterprise System" on page 128.](#))

To import a VM from an everRun MX system, first use XenConvert 2.1 to export OVF and Virtual Hard Disk (VHD) files from the everRun MX system to a network share, and then use the everRun Availability Console to import the OVF and VHD files from the network share to the everRun Enterprise system.



Caution: Consider backing up your source VM before preparing it for export from the everRun MX system.

Notes:

- You can import only VMs running Windows 7 or Windows Server 2008 from everRun MX systems.
- For Windows-based VMs, you must install VirtIO drivers in the guest operating system **before** exporting the VM from the everRun MX system, as described in this topic. If you do not install the VirtIO drivers, the imported VM crashes while booting on the everRun Enterprise system.
- You need to map a network share that is accessible from the source VM on the everRun MX system and also accessible by the management PC running everRun Availability Console. You use XenConvert to export the VM to this share, then you import the VM to the everRun Enterprise system from this share.
- Your source VM must be shut down while you are exporting the OVF file from the everRun MX system. Consider scheduling a planned maintenance period for this process.
- The time required for the export and import 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 1 Gb/s network may take about 30 minutes each way, export and import.
- If you will continue to use the source VM on the everRun MX system after the import, remember to set a different MAC address and IP address for the VM on the everRun Enterprise system.
- If the everRun Enterprise system switches from the primary PM to the secondary PM during an import, the import process fails. This does not affect the continuous uptime of your system, but you must delete the incomplete VM and associated volumes on the everRun Enterprise system, and import them again.



Exporting an OVF File from the everRun MX System

Exporting a VM from the everRun MX system exports the VM's configuration in an OVF file along with a copy of the selected volumes on your management PC.

To prepare for exporting a VM from the everRun MX system

1. Log on to the everRun MX system with **Citrix XenCenter**.
2. In the left-hand navigation panel, select the VM to export.
3. Click the **Console** tab 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](#)" on page 151.
5. Run the Windows System Preparation Tool (`Sysprep`) to prepare the guest operating system for redeployment.
6. Install the VirtIO drivers and the XenConvert utility in the Windows guest operating system:
 - a. Download the **VirtIO.exe** driver installation utility from the **Drivers and Tools** section of the **everRun Enterprise Downloads and Support** page at <http://www.stratus.com/go/support/everrun> to the guest operating system. This installation utility installs the VirtIO drivers and also the XenConvert utility required for an export from the everRun MX system.
 - b. Right-click the installation utility and click **Run as administrator**.
 - c. Click **OK** to install the software, and monitor the progress in the command prompt window.
 - d. Click **Restart Later** when Windows prompts you to restart the guest operating system.



Note: Windows prompts you to restart while the installation utility is still working. **Do not restart the VM** until you complete the following steps; otherwise, the driver installation fails and your imported VM will not boot on the everRun Enterprise system.

- e. Wait until the command prompt window indicates that the installation is finished and prompts you to **Press any key to continue**.
- f. Click the command prompt window to make it the active window, then press any key and wait for command prompt window and WinZip window to close.
- g. Restart the guest operating system to load the new drivers.

You can optionally uninstall the VirtIO drivers and the XenConvert utility after a successful import, as described later in this topic.

To export the VM and boot volume from the everRun MX system

1. In the Windows guest operating system on the everRun MX system, map a network share to which you will export the VM. For example, you can access a network share on your management PC that runs the everRun Availability Console.
2. Start **Citrix XenConvert** in the source VM.
3. Verify that **From: This machine** is selected.
4. Select **To: Open Virtualization Format (OVF) Package**. Click **Next**.
5. Select only the **(Boot)** volume to export. Explicitly deselect other volumes by clicking the **Source Volume** pulldown menu and selecting **None**. Do not change any other settings on this page. Click **Next**.



Note: You can export only one volume at a time; otherwise, the export fails. See the next procedure to export additional volumes.

6. Specify a path in the **Please choose a folder to store the Open Virtualization (OVF) package** text area. Click **Browse** and select a new, empty folder on the network share that you mounted for the export.
7. Ensure that the following XenConvert options are disabled. These options are not supported, and they can prevent a successful import:
 - Include a EULA in the OVF package
 - Create Open Virtual Appliance (OVA)
 - Compress Open Virtual Appliance (OVA)
 - Encrypt
 - Sign with Certificate
8. Click **Next**.
9. Optionally modify the name of the target OVF file. Click **Next**.
10. Click **Convert**.



Note: During the export process, if Windows displays a message indicating that you need to format a hard disk to use it, you can click **Cancel** to dismiss this message. The export continues normally.

To export each additional volume from the VM on the everRun MX system

1. Restart **Citrix XenConvert** on the source VM.
2. Verify that **From: This machine** is selected.
3. Select **To: XenServer Virtual Hard Disk (VHD)**. Click **Next**.
4. Select only **one** volume to export. Explicitly deselect other volumes by clicking the **Source Volume** pulldown menu and selecting **None**.

Do not change any other settings on this page. Click **Next**

5. Specify a path in the **Please choose a folder to store the Open Virtualization (OVF) package** text area. Click **Browse** and select a new, empty folder on the network share that you mounted for the export. Click **Next**.



Note: XenConvert does not give the option of specifying VHD file names, so each VHD export must initially be stored in a different folder to avoid overwriting the previous files.

6. Click **Convert**. This creates a VHD file and a PVP file.
7. After the VHD export, rename the new VHD file to give it a new, unique name and move it to the folder with the boot volume OVF and VHD. The PVP file is not used.
8. Repeat this procedure for each additional volume.

Importing the OVF file to the everRun Enterprise System

Importing a VM to the everRun Enterprise system imports the VM's configuration and any associated volumes that you select from your exported files.

Prerequisites:



- The selected OVF file (boot volume) and all associated VHD files (additional volumes) must be in the same directory, and no other VHD files can be in that directory.
- Both PMs of the everRun Enterprise system must be online for the import process to function properly.

To import a VM to the everRun Enterprise system

1. If applicable, on your management PC, map the network share containing the exported OVF and VHD files.
2. Log on to the everRun Enterprise system with the everRun Availability Console.
3. On the **Physical Machines** page (see ["The Physical Machines Page" on page 80](#)), verify that both PMs are in the **running** state and that neither PM is in maintenance mode or in the process of synchronizing.
4. On the **Virtual Machines** page (see ["The Virtual Machines Page" on page 84](#)), click **Import/Restore** to open the import wizard.
5. If prompted, allow the required Java™ plugins to load in your web browser. For information about enabling Java for the everRun Availability Console, see ["Compatible Internet Browsers" on page 29](#).
6. Click **Browse**. In the file browser, select the **.ovf** file that you want to import from your management PC and click **Import**.
7. Review the information and make any desired edits, if necessary:
 - **Name, CPU, and Memory**

Change the name of the virtual machine, edit the number of vCPUs, or allocate the total memory it can use.
 - **Storage**

Shows all of the volumes. Select the **Create** box for a volume to allocate a storage container for the volume on the everRun Enterprise system (the boot volume is required). Select the **Restore Data** box to import data for a volume from the OVF 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. If you will continue running the source VM on the everRun MX system, ensure that the MAC address for each network interface in the new VM is different from the source VM.

8. 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 on the everRun Enterprise system.
9. Click **Import** to begin importing the VM. When the transfer is complete, click **Done** to close the import wizard.



Note: Imported volumes begin to appear on the **Volumes** page of the everRun Availability 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.

10. If applicable, use the **Reprovision Virtual Machine** wizard to allocate additional resources to the VM, as described in ["Reprovisioning Virtual Machine Resources" on page 161](#).

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

11. Click **Console** to open the console of the VM and log on to the guest operating system.
12. For Windows-based VMs, allow the guest operating system to automatically install the VirtIO drivers and other required drivers, which may take several minutes. When a notification icon indicates that **Your device is ready to use** and you are prompted to restart, restart the guest operating system to load the drivers.
13. 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 everRun Enterprise system continues to synchronize data between PMs to enable High Availability (HA) or Fault Tolerant (FT) operation.



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

**To optionally uninstall the VirtIO drivers from the source VM on the everRun MX system
(Windows-based VMs only)**

After you successfully import the new VM to the everRun Enterprise system, you can uninstall the VirtIO drivers and the XenConvert utility from the Windows-based source VM on the everRun MX system. However, uninstalling this software is optional, as it does not interfere with the operation of the VM.

1. In the console of the source Windows-based VM, locate the **VirtIO.exe** installation utility.
(This same utility is used to uninstall the VirtIO drivers if they are present.)
2. Right-click the installation utility and click **Run as administrator**.
3. Click **OK** to uninstall the VirtIO drivers, and monitor the progress in the command prompt session.
4. When prompted, press any key to close the utility. No restart is necessary.

Troubleshooting

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

To clean up after a canceled or failed export from the everRun MX system

In the Windows guest operating system, consider saving log file information from XenConvert, then close the utility. Remove all of the files from the export folder on your network share, or create a new folder for a subsequent export. You must select an empty folder for each new export.

To clean up after a canceled or failed import on the everRun Enterprise system

In the everRun Availability Console, remove the imported VM and any volumes associated with the imported VM.

To recover from a failed export from the everRun MX system

The export fails if you export more than one volume at a time. Run XenConvert again and be careful to explicitly deselect all but one volume for the export. Also, ensure that you use an empty folder for each new export.

To recover from a failed import to the everRun Enterprise system

The imported VM crashes if the VirtIO drivers are not present in a Windows-based VM. Before running the XenConvert export again, ensure that you install the VirtIO drivers in the VM on the everRun MX system.

To recover missing data volumes in the VM on the everRun Enterprise system

If your data volumes do not appear in the VM on the everRun Enterprise 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.
- Use **Disk Management** to bring data volumes online.

To recover missing network devices in the VM on the everRun Enterprise system

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

Related Topics

["Migrating From Non-everRun Enterprise Systems" on page 97](#)

["Creating and Migrating Virtual Machines" on page 125](#)

["Configuring Windows-based Virtual Machines" on page 152](#)

["Configuring Linux-based Virtual Machines" on page 154](#)

["Managing Virtual Machine Resources" on page 160](#)

["Managing the Operation of a Virtual Machine" on page 156](#)

Importing an OVF File from an Avance System to the everRun Enterprise System

Import an Open Virtualization Format (OVF) file from an Avance unit if you want to transfer a VM to the everRun Enterprise system for deployment. (To migrate a physical machine (PM) or virtual machine (VM) to the everRun Enterprise system without using an OVF file, see ["Migrating a Physical Machine or Virtual Machine to the everRun Enterprise System" on page 128.](#))

To import a VM from an Avance unit, first use the Avance Management Console to export OVF and hard disk files to a management PC, and then use the everRun Availability Console to import the OVF and hard disk files from the management PC to the everRun Enterprise system.



Caution: Consider backing up your source VM before preparing it for export from the Avance unit.

Notes:

- You can import only VMs running CentOS/RHEL 6, Windows 7, Windows Server 2008, or Ubuntu 12.04 or later from Avance units.
- For Windows-based VMs, you must install VirtIO drivers in the guest operating system **before** exporting the VM from the Avance unit, as described in this topic. If you do not install the VirtIO drivers, the imported VM crashes while booting on the everRun Enterprise system.
- For Linux-based VMs, before exporting the VM from the Avance unit, consider editing the `/etc/fstab` file 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 everRun Enterprise system, your new VM may boot into single-user mode if it cannot mount the volumes with their original device names. You can restore the `/etc/fstab` entries in the new VM with the correct device names after the import process, as described below in **Troubleshooting**.
- For Ubuntu-based VMs, you must edit the `/boot/grub/grub.cfg` file and change the `gfxmode` parameter to `text` (for example, set `gfxmode=text`) before exporting the VM from the Avance unit; otherwise, the new VM's console hangs on the everRun Enterprise system. You can restore the original setting in the source VM after the migration.
- Your source VM must be shut down while you are exporting the OVF file or creating a snapshot on the Avance unit. Consider scheduling a planned maintenance period for this process.
- The time required for the export and import 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 1 Gb/s network may take about 30 minutes each way, export and import.
- If you will continue to use the source VM on the Avance unit after the import, remember to set a different MAC address and IP address for the VM on the everRun Enterprise system.





- If the everRun Enterprise system switches from the primary PM to the secondary PM during an import, the import process fails. This does not affect the continuous uptime of your system, but you must delete the incomplete VM and associated volumes on the everRun Enterprise system, and import them again.

Exporting an OVF File from the Avance Unit

Exporting a VM from the Avance unit exports the VM's configuration in an OVF file along with a copy of the selected volumes on your management PC.

To prepare for exporting a VM from the Avance unit (Windows-based VMs only)

1. Log on to the Avance unit with the Avance Management 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](#)" on page 151.
5. Run the Windows System Preparation Tool (`Sysprep`) to prepare the guest operating system for redeployment.
6. Install the VirtIO drivers in the Windows guest operating system:
 - a. Download the **VirtIO.exe** driver installation utility from the **Drivers and Tools** section of the **everRun Enterprise Downloads and Support** page at <http://www.stratus.com/go/support/everrun> to the guest operating system.
 - b. Right-click the installation utility and click **Run as administrator**.
 - c. Click **OK** to install the VirtIO drivers, and monitor the progress in the command prompt window.
 - d. Click **Restart Later** when Windows prompts you to restart the guest operating system.



Note: Windows prompts you to restart while the installation utility is still working. **Do not restart the VM** until you complete the following steps; otherwise, the driver installation fails and your imported VM will not boot on the everRun Enterprise system.

- e. Wait until the command prompt window indicates that the VirtIO driver installation is finished and prompts you to **Press any key to continue**.
- f. Click the command prompt window to make it the active window, then press any key and wait for command prompt window and WinZip window to close.
- g. Restart the guest operating system to load the new drivers.

Installing the VirtIO drivers also installs the XenConvert utility required for exports from everRun MX systems; however, this utility is not used on Avance units. You can optionally uninstall the VirtIO drivers and the XenConvert utility after a successful import, as described later in this topic.

To export a VM from the Avance unit

The following procedure describes how to export a VM from Avance, but you can also create a snapshot and export the snapshot to reduce downtime for the source VM. To create a snapshot, see the Avance online help.

1. Log on to the Avance unit with the Avance Management Console.
2. On the **Virtual Machines** page, select the VM to export.
3. With the VM selected, click **Shutdown** and wait for the VM to power off.
4. Click **Export** to display the export wizard.
5. If prompted, allow the required Java™ plugins to load in your web browser.
6. Click **Export VM**. (Click **Export Snapshot** if you created a snapshot.)
7. Click **Browse**. Select a location for the export on your management PC running the Avance Management Console and click **Save**.
8. Select the volumes you want capture, or click **VM Configuration Only** to include only the configuration details of each volume in the export file, but not the data.
9. Click **Export**.

Importing the OVF file to the everRun Enterprise System

Importing a VM to the everRun Enterprise system imports the VM's configuration and any associated volumes that you select from the OVF export on your management PC.



Prerequisite: Both PMs of the everRun Enterprise system must be online for the import process to function properly.

To import a VM to the everRun Enterprise system

1. Log on to the everRun Enterprise system with the everRun Availability Console.
2. On the **Physical Machines** page (see ["The Physical Machines Page" on page 80](#)), verify that both PMs are in the **running** state and that neither PM is in maintenance mode or in the process of synchronizing.
3. On the **Virtual Machines** page (see ["The Virtual Machines Page" on page 84](#)), click **Import/Restore** to open the import wizard.
4. If prompted, allow the required Java plugins to load in your web browser. For information about enabling Java for the everRun Availability Console, see ["Compatible Internet Browsers" on page 29](#).
5. Click **Browse**. In the file browser, select the **.ovf** file that you want to import from your management PC and click **Import**.
6. Review the information and make any desired edits, if necessary:
 - **Name, CPU, and Memory**

Change the name of the virtual machine, edit the number of vCPUs, or allocate the total memory it can use.
 - **Storage**

Shows all of the volumes. Select the **Create** box for a volume to allocate a storage container for the volume on the everRun Enterprise system (the boot volume is required). Select the **Restore Data** box to import data for a volume from the OVF 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. If you will continue

running the source VM on the Avance unit, ensure that the MAC address for each network interface in the new VM is different from the source VM

7. 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 on the everRun Enterprise system.
8. Click **Import** to begin importing the VM. When the transfer is complete, click **Done** to close the import wizard.



Note: Imported volumes begin to appear on the **Volumes** page of the everRun Availability 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.

9. If applicable, use the **Reprovision Virtual Machine** wizard to allocate additional resources to the VM, as described in ["Reprovisioning Virtual Machine Resources" on page 161](#).

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

10. Click **Console** to open the console of the VM and log on to the guest operating system.
11. For Windows-based VMs, allow the guest operating system to automatically install the VirtIO drivers and other required drivers, which may take several minutes. When a notification icon indicates that **Your device is ready to use** and you are prompted to restart, restart the guest operating system to load the drivers.
12. 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 everRun Enterprise system may continue to synchronize data between PMs to enable High Availability (HA) or Fault Tolerant (FT) operation.



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

To optionally uninstall the VirtIO drivers from the source VM on the Avance unit (Windows-based VMs only)

After you successfully import the new VM to the everRun Enterprise system, you can uninstall the VirtIO drivers and the XenConvert utility from the Windows-based source VM on the Avance unit.

However, uninstalling this software is optional, as it does not interfere with the operation or continuous uptime of the Avance unit.

1. In the console of the source Windows-based VM, locate the **VirtIO.exe** installation utility. (This same utility is used to uninstall the VirtIO drivers if they are present.)
2. Right-click the installation utility and click **Run as administrator**.
3. Click **OK** to uninstall the VirtIO drivers, and monitor the progress in the command prompt session.
4. When prompted, press any key to close the utility. No restart is necessary.

Troubleshooting

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

To clean up after a canceled or failed export from the Avance unit

On your management PC, remove all of the files from the export folder or create a new folder for a subsequent export.

To clean up after a canceled or failed import on the everRun Enterprise system

In the everRun Availability Console, remove the imported VM and any volumes associated with the imported VM.

To recover from a failed import to the everRun Enterprise system

The imported VM crashes if the VirtIO drivers are not present in a Windows-based VM. Before running the export again, ensure that you install the VirtIO drivers in the VM on the Avance unit.

To recover when the new VM's console hangs on the everRun Enterprise system

For Ubuntu-based VMs, the VM console hangs in everRun Availability Console if you do not properly set the `gfxmode` parameter before the import process (as described in **Notes**). If the VM console hangs, keep restarting the VM until the console opens properly in everRun Availability Console and then modify the `gfxmode` parameter to prevent subsequent issues.

For additional VM console troubleshooting, see ["Opening a Virtual Machine Console Session" on page 158](#).

To recover missing data volumes in the VM on the everRun Enterprise system

If your data volumes do not appear in the VM on the everRun Enterprise 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, from Avance (`/dev/xvda` through `/dev/xvdh`) to everRun Enterprise (`/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 everRun Enterprise system

If your network devices do not appear in the VM on the everRun Enterprise 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.

Related Topics

["Migrating From Non-everRun Enterprise Systems" on page 97](#)

["Creating and Migrating Virtual Machines" on page 125](#)

["Configuring Windows-based Virtual Machines" on page 152](#)

["Configuring Linux-based Virtual Machines" on page 154](#)

["Managing Virtual Machine Resources" on page 160](#)

["Managing the Operation of a Virtual Machine" on page 156](#)

Managing Windows Drive Labels

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



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, see:

<http://windows.microsoft.com/en-us/windows-vista/Change-add-or-remove-a-drive-letter>

Related Topics

["Creating and Migrating Virtual Machines" on page 125](#)

["Configuring Windows-based Virtual Machines" on page 152](#)

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:

- ["Creating and Initializing a Disk \(Windows-based VMs\)" on page 153](#)
- ["Installing Applications \(Windows-based VMs\)" on page 153](#)

In addition, ensure that you configure the following settings:

- Change the time zone in the VM to correspond to the time zone configured on the **Date and Time** preference page in everRun Availability Console (see ["Configuring Date and Time" on page 68](#)); otherwise, the VM's time zone changes whenever VMs restart or migrate. Network Time Protocol (NTP) is recommended for both the VM and the everRun Enterprise system.

- Disable hibernation (enabled by default in some cases) to prevent the guest operating system from going into a power-saving state.
- Configure the 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.

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 everRun Availability Console to create a new volume in a storage group on the everRun Enterprise system, as described in ["Creating a Volume in a Virtual Machine" on page 163](#).
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 everRun 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" on page 158](#)

["Configuring Windows-based Virtual Machines" on page 152](#)

["Creating and Migrating Virtual Machines" on page 125](#)

["Managing the Operation of a Virtual Machine" on page 156](#)

Installing Applications (Windows-based VMs)

Install an application in a Windows-based virtual machine over the network. For example, map a network share that contains the installation program or download the installation program to the guest operating system as an executable file or ISO file.



Note: You cannot use virtual CDs to install applications.

Related Topics

["Opening a Virtual Machine Console Session" on page 158](#)

["Configuring Windows-based Virtual Machines" on page 152](#)

["Creating and Migrating Virtual Machines" on page 125](#)

["Managing the Operation of a Virtual Machine" on page 156](#)

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 \(Linux-based VMs\)" on page 155](#)
- ["Installing Applications \(Linux-based VMs\)" on page 155](#)

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.
- Install the `kexec-tools` package and configure the guest operating system to generate a crash dump file if the system crashes.
- If the guest operating system is **Ubuntu Linux**:
 - To prevent a problem where the VM console hangs in everRun Availability 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 on first boot, keep restarting the VM until the console opens properly in everRun Availability Console and then modify the `gfxmode` parameter to prevent subsequent issues.

- For the minimal server version of Ubuntu, optionally install the `acpid` package to enable the **Shutdown** button in everRun Availability Console, as described in ["Shutting Down a Virtual Machine" on page 156](#).

For more information about these settings, see your Linux documentation.

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 everRun Availability Console, create a new volume in a storage group, as described in ["Creating a Volume in a Virtual Machine" on page 163](#).
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 everRun 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" on page 158](#)

["Configuring Linux-based Virtual Machines" on page 154](#)

["Creating and Migrating Virtual Machines" on page 125](#)

["Managing the Operation of a Virtual Machine" on page 156](#)

Installing Applications (Linux-based VMs)

Install an application in a Linux-based virtual machine over the network. For example, mount a network drive that contains the installation package or download the installation package to the guest operating system as an executable file or ISO file.



Note: You cannot use virtual CDs to install applications.

Related Topics

["Opening a Virtual Machine Console Session" on page 158](#)

["Configuring Linux-based Virtual Machines" on page 154](#)

["Creating and Migrating Virtual Machines" on page 125](#)

["Managing the Operation of a Virtual Machine" on page 156](#)

Managing the Operation of a Virtual Machine

Manage the operation of a virtual machine as described in:

- ["Starting a Virtual Machine" on page 156](#)
- ["Shutting Down a Virtual Machine" on page 156](#)
- ["Powering Off a Virtual Machine" on page 157](#)
- ["Opening a Virtual Machine Console Session" on page 158](#)
- ["Renaming a Virtual Machine" on page 159](#)
- ["Removing a Virtual Machine" on page 160](#)

For additional information configuration and troubleshooting information, see ["Advanced Topics \(Virtual Machines\)" on page 173](#).

Starting a Virtual Machine

Start a virtual machine to boot the guest operating system installed in the virtual machine.

To start a virtual machine

1. Select a virtual machine on the **Virtual Machines** page.
2. Click **Start**.

Related Topics

["Shutting Down a Virtual Machine" on page 156](#)

["Powering Off a Virtual Machine" on page 157](#)

["Managing the Operation of a Virtual Machine" on page 156](#)

Shutting Down a Virtual Machine

Shut down a virtual machine to begin an orderly shutdown of the guest operating system.



Note: You can shut down a virtual machine with guest operating system commands. Some guests allow (or can be configured to allow) you to shut down a virtual machine using the ever-Run Availability Console.

To shut down a virtual machine in everRun Availability Console

1. Select a virtual machine on the **Virtual Machines** page.
2. Click **Shutdown**.

If the virtual machine is not responding, you can also **Power Off** the virtual machine to stop it without properly shutting down the guest operating system.

Shutting down a virtual machine in the everRun Availability 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:

- 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 Linux, 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.

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 everRun Availability Console.

Related Topics

["Starting a Virtual Machine" on page 156](#)

["Powering Off a Virtual Machine" on page 157](#)

["Managing the Operation of a Virtual Machine" on page 156](#)

Powering Off a Virtual Machine

Power off a virtual machine 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 virtual machine is similar to pulling the power cord, which may result in data loss.

To power off a virtual machine

1. Select a virtual machine on the **Virtual Machines** page.
2. Click **Power Off**.

Related Topics

["Starting a Virtual Machine" on page 156](#)

["Shutting Down a Virtual Machine" on page 156](#)

["Managing the Operation of a Virtual Machine" on page 156](#)

["Advanced Topics \(Virtual Machines\)" on page 173](#)

Opening a Virtual Machine Console Session

Open a virtual machine (VM) console to display the console of the guest operating system running in the VM.

The following procedure describes how to open a VM console in the everRun Availability Console, but you can also use a remote desktop application for this purpose.

To open a VM console

1. Select a VM on the **Virtual Machines** page.
2. Ensure that the VM is in a running state.
3. Click **Console**.
4. If prompted, allow the required Java™ plugins to load in your web browser.

Troubleshooting

To resolve an issue where the VM console window does not open

Allow the required Java™ plugins to load in your web browser. For information about enabling Java for the everRun Availability Console, see ["Compatible Internet Browsers" on page 29](#).

If you still have problems opening the VM console session, you may need to 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 everRun Enterprise system

For Ubuntu-based VMs, the VM console hangs in everRun Availability Console if you do not properly set the `gfxmode` parameter. See ["Configuring Linux-based Virtual Machines" on page 154](#).

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 everRun Availability 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" on page 156](#)

["Shutting Down a Virtual Machine" on page 156](#)

["Managing the Operation of a Virtual Machine" on page 156](#)

Renaming a Virtual Machine

Rename a virtual machine 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 virtual machine, use guest operating system tools.

To rename a virtual machine

1. Locate the virtual machine on the **Virtual Machines** page.
2. Double-click the name of the virtual machine.
3. Specify the new name and press **Enter**.

Related Topics

["Removing a Virtual Machine" on page 160](#)

["Creating and Migrating Virtual Machines" on page 125](#)

["Managing the Operation of a Virtual Machine" on page 156](#)

Removing a Virtual Machine

Remove a virtual machine to permanently delete the virtual machine and optionally delete associated volumes from the everRun Enterprise system.

To remove a virtual machine

1. Select a virtual machine on the **Virtual Machines** page.
2. Click **Shutdown**.
3. When the virtual machine 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 virtual machine.



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 virtual machine and any selected volumes.

Related Topics

["Renaming a Virtual Machine" on page 159](#)

["Creating and Migrating Virtual Machines" on page 125](#)

["Managing the Operation of a Virtual Machine" on page 156](#)

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" on page 161](#)

To reconfigure virtual machine volumes, see the following task-specific topics:

- ["Creating a Volume in a Virtual Machine" on page 163](#)
- ["Attaching a Volume to a Virtual Machine" on page 164](#)
- ["Detaching a Volume from a Virtual Machine" on page 165](#)
- ["Removing a Volume from a Virtual Machine" on page 166](#)

To recover virtual machine resources, freeing space for new volumes or virtual CDs, see:

- ["Recovering Virtual Machine Resources" on page 168](#)

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" on page 120](#).
- 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" on page 84](#)).
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, Description, Protection and OS** page:
 - a. Type the **Name** and an optional **Description** for the VM as they will appear in the everRun Availability Console.

b. Select the level of protection to use for the VM:

- **High Availability (HA)** – Provides basic failover and recovery, with some faults requiring an (automatic) VM reboot for recovery. Use HA for applications that can tolerate some downtime and that do not need the downtime protection that FT provides.
- **Fault Tolerant (FT)** – Transparently protects an application by creating a redundant environment for a VM running across two physical machines. Use FT for applications that need greater downtime protection than HA provides.

For more information about these levels of protection, see ["Modes of Operation" on page 11](#).

c. Click **Next**.

5. On the **Volumes** page, you can:



Note: You cannot modify the VM boot volume, only data volumes.

- 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.) Specify the **Name**, **Size**, and **Storage Group** for the new volume.
- Click **Detach** to disconnect a volume from a VM and keep it for future use.
- Click **Delete** to permanently remove a volume from the everRun system.
- Select an unattached volume from a pulldown menu (if displayed) and click **Attach**.

For more information, see ["Planning Virtual Machine Storage" on page 123](#). To continue, click **Next**.

6. On the **Networks** page, select the shared networks to assign to the VM. For more information, see ["Planning Virtual Machine Networks" on page 124](#). To continue, click **Next**.

7. On the **vCPUs and Memory** page, specify the number of **vCPUs** and the amount of **Memory** to assign to the VM. For more information, see ["Planning Virtual Machine vCPUs" on page 120](#) and ["Planning Virtual Machine Memory" on page 122](#). 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" on page 160](#)

["Planning Virtual Machine Resources" on page 120](#)

["Managing Virtual Machines" on page 119](#)

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" on page 164](#).)



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" on page 84](#)).
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** to skip the **Name, Description, Protection and OS** page. (If applicable, see ["Reprovisioning Virtual Machine Resources" on page 161](#) 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. Specify the **Name**, **Size**, and **Storage Group** for the new volume.
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 Windows or Linux guest operating system, as described in:
 - ["Creating and Initializing a Disk \(Windows-based VMs\)" on page 153](#)
 - ["Creating and Initializing a Disk \(Linux-based VMs\)" on page 155](#)

Related Topics

["Detaching a Volume from a Virtual Machine" on page 165](#)

["Removing a Volume from a Virtual Machine" on page 166](#)

["Managing Virtual Machine Resources" on page 160](#)

["Planning Virtual Machine Resources" on page 120](#)

["Managing Virtual Machines" on page 119](#)

Attaching a Volume to a Virtual Machine

Attach a volume to connect a currently unused volume to a virtual machine.



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 **VM** column is **None**.
2. Open the **Virtual Machines** page (see ["The Virtual Machines Page" on page 84](#)).
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** to skip the **Name, Description, Protection and OS** page. (If applicable, see ["Reprovisioning Virtual Machine Resources" on page 161](#) 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 everRun Enterprise 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" on page 163](#)

["Detaching a Volume from a Virtual Machine" on page 165](#)

["Removing a Volume from a Virtual Machine" on page 166](#)

["Managing Virtual Machine Resources" on page 160](#)

["Planning Virtual Machine Resources" on page 120](#)

["Managing Virtual Machines" on page 119](#)

Detaching a Volume from a Virtual Machine

Detach a volume to disconnect it from a virtual machine and keep it for future use. (You can also permanently delete it from the everRun Enterprise system, as described in ["Removing a Volume from a Virtual Machine" on page 166](#).)



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" on page 84](#)).
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** to skip the **Name, Description, Protection and OS** page. (If applicable, see ["Reprovisioning Virtual Machine Resources" on page 161](#) 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. Above the volume name, click **Detach** 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" on page 164](#)

["Removing a Volume from a Virtual Machine" on page 166](#)

["Managing Virtual Machine Resources" on page 160](#)

["Planning Virtual Machine Resources" on page 120](#)

["Managing Virtual Machines" on page 119](#)

Removing a Volume from a Virtual Machine

Remove a volume to permanently delete it from the everRun Enterprise system. (You can also detach a volume from the virtual machine but keep it for future use, as described in ["Detaching a Volume from a Virtual Machine" on page 165](#).)



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" on page 84](#)).
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** to skip the **Name, Description, Protection and OS** page. (If applicable, see ["Reprovisioning Virtual Machine Resources" on page 161](#) 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. Above the volume name, click **Delete** 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 **VM** column must read **None**, otherwise, the **Remove** button is not displayed.)
3. Click **Remove**.

Related Topics

["Detaching a Volume from a Virtual Machine" on page 165](#)

["Attaching a Volume to a Virtual Machine" on page 164](#)

["Managing Virtual Machine Resources" on page 160](#)

["Planning Virtual Machine Resources" on page 120](#)

["Managing Virtual Machines" on page 119](#)

Renaming a Volume on the everRun System

Rename a volume on the everRun Enterprise 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 everRun 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" on page 163](#)

["Detaching a Volume from a Virtual Machine" on page 165](#)

["Removing a Volume from a Virtual Machine" on page 166](#)

["Managing Virtual Machine Resources" on page 160](#)

["Planning Virtual Machine Resources" on page 120](#)

["Managing Virtual Machines" on page 119](#)

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" on page 160](#)
- ["Removing a Volume from a Virtual Machine" on page 166](#)
- ["Removing a Virtual CD" on page 172](#)

Related Topics

["Managing Virtual Machine Resources" on page 160](#)

["Planning Virtual Machine Resources" on page 120](#)

["Managing Virtual Machines" on page 119](#)

Managing Virtual CDs

Create and manage virtual CDs (VCDs) to make software installation media available to the virtual machines on your everRun Enterprise system in ISO format.

A VCD is a read-only ISO image file that resides on a storage device of the everRun system. Use the **Virtual CD Creation Wizard** (in everRun Availability Console) to upload an existing ISO file or to create a new ISO file from a physical CD/DVD source, as described in ["Creating a Virtual CD" on page 169](#).

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.



Note: To ensure continuous uptime, the everRun software prevents you from inserting a VCD after guest installation because it would prevent the system from migrating VMs in the event of a failure. However, you can still boot a virtual machine from a VCD for troubleshooting purposes.

You manage VCDs as described in:

- ["Creating a Virtual CD" on page 169](#)
- ["Burning a CD or DVD for a Virtual CD" on page 171](#)
- ["Booting from a Virtual CD" on page 171](#)
- ["Renaming a Virtual CD" on page 172](#)
- ["Removing a Virtual CD" on page 172](#)

Creating a Virtual CD

Create a virtual CD (VCD) to make software installation media available to the virtual machines (VM) on your everRun system.

To create a VCD, use the **Virtual CD Creation Wizard** to copy an ISO file to a storage device on the everRun system. Thereafter, you can boot from it (see ["Booting from a Virtual CD" on page 171](#)) to install a guest operating system or start a VM from a bootable recovery VCD.

Notes:



1. You cannot use VCDs to install applications in virtual machine. If necessary, mount a network drive or ISO image in the guest operating system to access application media.
2. Each VCD consumes disk space in the storage group in which it is stored. Unless you use a VCD on a regular basis, remove it when it is no longer needed.
3. 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 everRun Availability Console.
3. Click **Create VCD** to open the **Virtual CD Creation Wizard**.
4. In the wizard, select a storage group with sufficient free space for the VCD.
5. Type a name for the VCD.
6. Select a source for the VCD:
 - **Upload ISO file** uploads a file from your remote system running the everRun Availability Console.
 - **Copy CD ISO from network source** copies the file from a Web URL.
7. If you selected **Upload ISO file**, click **Next** and select the ISO file to upload.
8. Click **Finish** to upload or copy the ISO file from the specified source.

The **Virtual CD Creation Wizard** reports that the VCD has been successfully added, but transferring the image may take several more minutes depending on the size of the image.

You can determine the status of a VCD by checking the status 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 use.

Related Topics

["Burning a CD or DVD for a Virtual CD" on page 171](#)

["Managing Virtual CDs" on page 169](#)

["Creating and Migrating Virtual Machines" on page 125](#)

Burning a CD or DVD for a Virtual CD

If you need to burn a physical CD or DVD that you will subsequently use to create a virtual CD on the ever-Run Enterprise system (see ["Creating a Virtual CD" on page 169](#)), note the following guidelines:

- Use only media-burning software and CD-R or DVD-R media and drives that support the Disk At Once (DAO) method. Using software that defaults to DAO mode, like ImgBurn from <http://imgburn.com/>, provides more assurance that you will use DAO mode.
- Always use new media.
- To minimize the probability of having buffer under-runs during the burn process, download the ISO image to the same computer on which you are burning the media.
- Always verify the newly burned disc. You can use the verification feature of your media-burning software.

Related Topics

["Creating a Virtual CD" on page 169](#)

["Managing Virtual CDs" on page 169](#)

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 from a virtual machine from a VCD

1. If necessary, create a VCD from a bootable CD/DVD (see ["Creating a Virtual CD" on page 169](#)).
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**.
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" on page 169](#)

["Managing Virtual CDs" on page 169](#)

["Creating and Migrating Virtual Machines" on page 125](#)

["Managing the Operation of a Virtual Machine" on page 156](#)

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" on page 172](#)

["Creating a Virtual CD" on page 169](#)

["Managing Virtual CDs" on page 169](#)

Removing a Virtual CD

Remove a virtual CD (VCD) to permanently delete it from the everRun Enterprise system.

To remove a VCD

1. In the everRun Availability Console, click **Virtual CDs**.
2. Locate the VCD you want to remove in the list.
3. Ensure that the **Can Remove** value for the VCD is **Yes**. If the value is **No**, the VCD is currently in use.
4. Select the VCD and click **Remove**.

Related Topics

["Renaming a Virtual CD" on page 172](#)

["Creating a Virtual CD" on page 169](#)

["Managing Virtual CDs" on page 169](#)

Advanced Topics (Virtual Machines)

The following topics describe procedures and information for advanced users:

- ["Assigning a Specific MAC Address to a Virtual Machine" on page 173](#)
- ["Selecting a Preferred PM for a Virtual Machine" on page 174](#)
- ["Changing the Protection Level for a Virtual Machine \(HA or FT\)" on page 174](#)
- ["Configuring the Boot Sequence for Virtual Machines" on page 175](#)
- ["Resetting MTBF for a Failed Virtual Machine" on page 176](#)
- ["Locating a Dump File in a Virtual Machine" on page 177](#)

To manage the operation of a virtual machine, see ["Managing the Operation of a Virtual Machine" on page 156](#).

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.



Note: The everRun Enterprise 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).



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" on page 84](#)).
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" on page 161](#) 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 everRun software to automatically assign the MAC address.
7. Click **Finish**.

Related Topics

["Advanced Topics \(Virtual Machines\)" on page 173](#)

["Managing Virtual Machine Resources" on page 160](#)

["Managing the Operation of a Virtual Machine" on page 156](#)

Selecting a Preferred PM for a Virtual Machine

Select a preferred physical machine to ensure that a virtual machine runs on a particular physical machine in the everRun Enterprise 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 \(Virtual Machines\)" on page 173](#)

["Managing the Operation of a Virtual Machine" on page 156](#)

Changing the Protection Level for a Virtual Machine (HA or FT)

You can change the protection level of guest VMs from high availability (HA) to fault tolerance (FT), or vice versa.

To change the protection level

1. On the **Virtual Machines** page, select a stopped VM (marked "stopped" in the **Activity** column).
(See ["Shutting Down a Virtual Machine" on page 156](#) for information about stopping a VM.)
2. In the bottom pane, click **Config** to open the **Reprovision Virtual Machine** wizard
3. On the **Configuring CPU and Memory** page, select the **HA** or **FT** button.
4. Press **Finish** and then **OK** (if the reconfiguration was successful).

Related Topics

["Modes of Operation" on page 11](#) (HA or FT)

["Advanced Topics \(Virtual Machines\)" on page 173](#)

["Managing the Operation of a Virtual Machine" on page 156](#)

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 everRun Enterprise 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 everRun 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 everRun 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 everRun boots the DNS VM.
 - 2 2 minutes after the DNS VM is started, everRun starts the App and DB servers in group 2.
 - 3 10 minutes after the DB VM is started, everRun 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 \(Virtual Machines\)" on page 173](#)

["Managing the Operation of a Virtual Machine" on page 156](#)

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 crashes, everRun Enterprise automatically restarts it, unless it has fallen below its MTBF threshold. If the virtual machine is below the MTBF threshold, everRun 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.



Note: The **Reset Device** button is displayed only if the virtual 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**.

Related Topics

["Advanced Topics \(Virtual Machines\)" on page 173](#)

["Managing the Operation of a Virtual Machine" on page 156](#)

["Creating a Diagnostic File" on page 71](#)

Locating a Dump File in a Virtual Machine

Locate a dump file in a virtual machine (VM) if the VM has crashed and you need to collect the dump file for troubleshooting purposes.

To collect a dump file for your service representative

- For Windows-based VMs—Retrieve the file from `C:\WINDOWS\MEMORY.DMP` (by default) in the file system of the VM.
- For Linux-based VMs—Retrieve the dump file from the `/var/crash` directory (by default) in the file system of the VM.

If you cannot locate a dump file, ensure that you configured the guest operating system to generate a crash dump file when the operating system hangs:

- Windows-based VMs: 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.
- Linux-based VMs: Install the `kexec-tools` package and enable crash dumps. For more information, see your Linux documentation.

Related Topics

["Advanced Topics \(Virtual Machines\)" on page 173](#)

["Managing the Operation of a Virtual Machine" on page 156](#)

["Creating a Diagnostic File" on page 71](#)

9

Chapter 9: Maintaining Physical Machines

You can maintain PMs in an everRun Enterprise system by adding or replacing various components or even the entire PM.



Prerequisite: Before you add, replace, or upgrade a component, see ["Physical Machine Hardware Maintenance Restrictions" on page 179](#).

Determine what component needs to be replaced and then read the topic for the appropriate procedure:

- To add or replace PM components, see:
 - ["Adding or Replacing Hot-Swappable Components" on page 180](#) for hot-swappable components such as network cables, fans, and power supplies
 - ["Adding or Replacing Components That Are Not Hot-Swappable" on page 181](#) for components such as CPUs, memory, and NICs.
- To replace a PM or a failed motherboard or RAID controller, see ["Replacing Physical Machines, Motherboards, NICs, or RAID Controllers" on page 182](#).
- To upgrade both PMs in a running system, see ["Upgrading Both Physical Machines In a Running System" on page 185](#).

For information on disks, see ["Logical Disks and Physical Disks" on page 15](#).

Physical Machine Hardware Maintenance Restrictions

When you replace physical machines (PMs), motherboards, or RAID controllers, you should ensure compatibility by complying with these restrictions:

- New PMs must have processors that are from the same processor family as the existing PM, in order to support live migration. If the processors in the new and existing PMs are from different processor families, you must stop the VMs to migrate them from the existing PM to the new PM.
- CPUs on a replacement PM must be compatible with the CPUs on the original PM.
- In the replacement PM, the quantity of the following resources must be the same or greater than in the original PM:
 - Number of processor cores.
 - Total memory.
 - Total logical disk capacity.
 - Total number of network ports; each port must support, at a minimum, the speed of the existing ports, and all add-on NICs within a particular PM must have the same vendor/model number.
 - Total number of network ports; each port must support, at a minimum, the speed of the existing ports.

In addition, check ["System Requirements Overview" on page 22](#) for information about system hardware and software requirements before performing hardware maintenance on a PM, to confirm that the maintenance you are planning complies with any system restrictions.

Related Topics

["Maintenance Mode" on page 109](#)

["Maintaining Physical Machines" on page 179](#)

["The everRun Availability Console" on page 51](#)

["Physical Machines and Virtual Machines" on page 8](#)

["The Physical Machines Page" on page 80](#)

Adding or Replacing Hot-Swappable Components

Use this procedure to add or replace a component that is hot-swappable. Such components may include network cables, fans, and power supplies. In this procedure, the PM continues to run.



Prerequisite: Before you add, replace, or upgrade a component, see ["Physical Machine Hardware Maintenance Restrictions" on page 179](#).

To add or replace a hot-swappable component

1. Determine which PM (node0 or node1) requires the component.
2. In the everRun Availability Console, click **Physical Machines** in the left-hand navigation panel.
3. 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)**.
4. Follow vendor instructions for adding or replacing a hot-swappable component in the PM.
5. Select the repaired PM on the **Physical Machines** page. Click **Finalize** and then click **OK**.

If you add a cable to both PMs and if they are on the same subnet, everRun detects the connectivity and pairs the NICs in a newly created shared network. You have the option of renaming the newly created shared network on the **network** page.

Related Topics

["Maintenance Mode" on page 109](#)

["Maintaining Physical Machines" on page 179](#)

["The everRun Availability Console" on page 51](#)

["Physical Machines and Virtual Machines" on page 8](#)

["The Physical Machines Page" on page 80](#)

Adding or Replacing Components That Are Not Hot-Swappable

Use this procedure to add or replace a component that is not hot-swappable. Such components may include CPUs and memory as well as fans and power supplies that are not hot-swappable.

In this procedure, you gracefully shut down a running PM.



Prerequisite: Before you add, replace, or upgrade a component, read ["Physical Machine Hardware Maintenance Restrictions" on page 179](#).

To add or replace a component that is not hot-swappable

1. Determine which PM (node0 or node1) or if each PM requires the replacement component.
2. In the everRun Availability Console, click **Physical Machines** in the left-hand navigation panel.
3. 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).

4. After the PM displays **running (in Maintenance)**, click **Shutdown** and then **OK**.
5. Add or replace the component.
6. If you disconnected any network cables, reconnect them. Do not add cables to any new network ports in this step.
7. On the PM that is shutdown, press the power button. As the PM powers on, everRun also powers on and begins synchronizing the PM's storage (🔄 appears).
8. On the **Networks** page, click the **Fix** button, if it is highlighted, which may occur when network cables have been moved on the upgraded PM.
9. Select the repaired PM on the **Physical Machines** page. Click **Finalize** and then click **OK**.
10. When synchronization is complete (🔄 disappears), perform steps 3 through 9 for the other PM, if necessary.



Note: To avoid data loss, do not power down the primary PM while the disks are synchronizing.

Related Topics

["Maintenance Mode" on page 109](#)

["Maintaining Physical Machines" on page 179](#)

["The everRun Availability Console" on page 51](#)

["Physical Machines and Virtual Machines" on page 8](#)

["The Physical Machines Page" on page 80](#)

Replacing Physical Machines, Motherboards, NICs, or RAID Controllers

You replace motherboards, NICs, RAID controllers, and a physical machine (PM) in an everRun Enterprise system while the system is running. You can replace PMs to upgrade a PM or to replace a failed PM. You can replace motherboards, NICs, or RAID controllers. Several types of hardware faults can hang or crash a PM in an everRun Enterprise system, including a failure of the motherboard, CPU, mid-plane, or storage controller. (If you want to recover a failed PM instead of replacing it, see ["Recovering a Failed Physical Machine" on page 114.](#))

You replace the PM or a component using vendor instructions. First, read "[Physical Machine Hardware Maintenance Restrictions](#)" on page 179 and then replace the hardware using the everRun PM **Replace** function, which deletes a PM from the everRun system's database, waits for the repaired/replaced PM to power on, and then images the repaired PM and synchronizes its storage from the running primary PM.

Prerequisite: Obtain installation software for the everRun Enterprise release that the PM has been running using one of the following methods:

- Download an install ISO from your authorized Stratus service representative.
- Extract an install ISO into the current working directory from the most recently used upgrade kit by executing a command similar to the following (*x.x.x.x* is the release number and *nnn* is the build number):

```
tar -xzf everRun_Enterprise_upgrade-x.x.x.x-nnn.kit *.iso
```

After you obtain the correct install ISO, save it or burn it to a DVD. See "[Obtaining everRun Enterprise Software](#)" on page 34.

Prerequisites: If you are replacing a PM, prepare the new PM:

1. Configure networks. See "[Network Architecture Overview](#)" on page 16.
2. Configure storage. See "[Storage Requirements](#)" on page 23.
3. Connect power. See "[Connecting Power](#)" on page 33.
4. Configure BIOS. See "[Configuring the BIOS](#)" on page 35.

To replace a failed PM, motherboard, NIC, or RAID controller

1. In the everRun Availability 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 shows **running (in Maintenance)**, click **Replace** (.

The everRun Enterprise software deletes the PM from the everRun system's database, and then prompts you to replace the PM or component.

4. Manually power off the old PM.

5. Install the new PM or component. If you are replacing a motherboard, NIC, or RAID controller, do so now. If you are replacing the PM, disconnect and remove it now, and then install the new PM. Check that a monitor and keyboard are connected to it.
6. Reconnect any network cables to their original positions. Check that Ethernet cables are connected from the new PM (or new NIC) to the network or directly to the running (primary) PM, if the two everRun system PMs are in close proximity. One Ethernet cable should connect from the first embedded port on the new PM or from a NIC port if the new PM does not have an embedded port.
7. Manually power on the PM. As the PM powers on, enter the BIOS and set the Optical Drive as the first boot device.
8. Either mount the ISO image or insert the DVD into the PM.
9. At the **Welcome** screen, select **Replace PM, Join system: Initialize data** and press **Enter**.



Note: If necessary, see ["Installing Software on the Second PM" on page 43](#) for reference. Although that topic refers to the "second PM," in this instance, it applies to the replacement PM.

10. When prompted, respond to the **Select interface for private Physical Machine connection**, and then respond to the prompt **Select interface for managing the system (ibiz0)**.
11. When prompted to configure **ibiz0**, select **Automatic configuration via DHCP** or **Manual Configuration (Static Address)**. (The installation software configures priv0 automatically.)
12. When installation is complete, the PM ejects the install DVD (if used) and reboots.
13. As the PM boots up, the everRun Availability Console displays its activity on the **Physical Machines** page. The **Activity** column displays the new PM as **recovery (in Maintenance)** and then **running** after recovery is complete.

Related Topics

["Maintenance Mode" on page 109](#)

["Maintaining Physical Machines" on page 179](#)

["The everRun Availability Console" on page 51](#)

["Physical Machines and Virtual Machines" on page 8](#)

["The Physical Machines Page" on page 80](#)

Upgrading Both Physical Machines In a Running System



Prerequisite: Before you upgrade to new physical machines, see ["Physical Machine Hardware Maintenance Restrictions" on page 179](#).

To upgrade to new physical machines

1. Upgrade everRun Enterprise software if required to support the new PM. See the everRun **Release Notes** and the help on the **everRun Upgrade Kits** page of the everRun Availability Console
2. Upgrade the first PM; see ["Replacing Physical Machines, Motherboards, NICs, or RAID Controllers" on page 182](#).
3. Repeat for the second PM. The everRun software then migrates the VMs to the other PM.
4. If you added additional NIC ports, see ["Network Architecture Overview" on page 16](#).

Related Topics

["Maintenance Mode" on page 109](#)

["Maintaining Physical Machines" on page 179](#)

["The everRun Availability Console" on page 51](#)

["Physical Machines and Virtual Machines" on page 8](#)

["The Physical Machines Page" on page 80](#)

Part 2: Supporting Documents

See the following support documents for release information, and reference and troubleshooting information.

- ["everRun Enterprise Release 7.1.0.0 Release Notes" on page 188](#)
- ["everRun Enterprise Command Line Interface Reference" on page 194](#)
- ["System Reference Information" on page 316](#)
- ["SNMP" on page 324](#)

10

Chapter 10: everRun Enterprise Release 7.1.0.0 Release Notes

These release notes are for everRun Enterprise release 7.1.0.0. See the following sections:

- [Important Considerations](#)
- [Known Issues](#)
- [New Features, Enhancements, and Bug Fixes](#)
- [Getting Help](#)



Note: For the latest technical information and updates, see the English version of the everRun Enterprise User's Guide at the [everRun Enterprise Downloads and Support](#) page at <http://www.stratus.com/go/support/everrun>.

Important Considerations

Upgrading From Previous Releases of everRun Enterprise

You can upgrade from everRun Enterprise release 7.0.x to everRun Enterprise release 7.1.x **with no VM downtime**, by following the instructions in ["Upgrading everRun Enterprise Software" on page 95](#).

If you are upgrading from any other release (for example everRun Enterprise 7.1.0.0 beta release) you must perform a complete system re-install.



Caution: All PMs and VMs must be in good health prior to performing an upgrade of everRun Enterprise software. Before starting an upgrade, examine the everRun Availability Console to verify that there are no alerts indicating PM or VM problems.



Note: The everRun Availability Console does not automatically refresh after an everRun Enterprise upgrade completes, giving the impression that the upgrade has not completed. Periodically refresh the everRun Availability Console during the upgrade to see if the upgrade has completed. To do this, click the browser's reload or refresh button. Pressing the F5 key also accomplishes this on many browsers.

If you encounter problems during an upgrade, contact Stratus support for assistance at <http://www.stratus.com/go/support/everrun>.

Guest Volumes Must be Smaller Than 2 TB

Guest volumes equal to or greater than 2 TB are not supported in this release. Guest volumes must be less than 2 TB.

The Active Directory Feature is Not Supported

The active directory feature is not supported in this release and has been removed from the everRun Availability Console. Some localized versions of the everRun Enterprise User's Guide prepared prior to the release of 7.1.0.0 may contain references to the active directory feature. For the latest information, see the English version of the everRun Enterprise User's Guide at the **everRun Enterprise Downloads and Support** page at <http://www.stratus.com/go/support/everrun>.

Optimizing Performance of A-Link Networks

Stratus recommends that you enable jumbo frames on A-Link networks by setting their Ethernet frame MTU size to 9000 bytes (by default, they are set to 1500 bytes). Doing so improves VM performance and reduces host processing overhead.

The A-Link networks must:

- Be a single Ethernet cable, point-to-point connection, or
- Have intermediate components (for example, switches) that are all fully capable of passing jumbo frame traffic.

Use AVCLI commands to enable jumbo frames. AVCLI is installed on the host system along with everRun Enterprise software. You can run AVCLI by logging onto the host through a remote console using the system IP address. Alternatively, you can install AVCLI on a remote management computer. For information about how to install AVCLI on a remote computer see ["AVCLI Command Overview" on page 194](#).

To set A-Links to use jumbo frames:

1. From a remote console computer, issue the `network-info` command to determine the name(s) of the A-Link network(s). In the command's output, look for the names of networks that have **role = A-Link**. See ["network-info" on page 257](#) for an example.
2. Issue the ["network-change-mtu" on page 255](#) command to change the MTU size to the maximum value of 9000 bytes. The changes takes effect immediately. The following example changes the `sync_2003` and `sync_2004` A-Link networks to use jumbo frames.

```
avcli network-change-mtu sync_2003 sync_2004 9000
```

3. Issue the `network-info` command to verify that the A-Links now have MTU values of 9000.



Caution: After you issue the `network-change-mtu` command, do not issue another `network-change-mtu` command until after the new mtu settings have taken affect. Use the `network-info` command as described in step 3 above, to verify that the new mtu settings have taken affect.

Migrating a PM or VM to an everRun Enterprise System

Migrating Windows 2012 R2 or Windows 8.x PMs or VMs from a non-everRun Enterprise system to an everRun Enterprise system is not supported. For a list of PM and VM operating systems that you can migrate, see ["Migrating a Physical Machine or Virtual Machine to the everRun Enterprise System" on page 128](#).

Status of RAID Physical Disks is Not Monitored

everRun Enterprise software does not monitor the state of physical disks in a RAID set. You must use the tools supplied by the RAID controller vendor to monitor the health and status of individual physical disks in a RAID set.

Other Important everRun Enterprise Considerations

For important considerations about everRun Enterprise systems, see ["Important Physical Machine and Virtual Machine Considerations" on page 319](#).

Known Issues

Replacing Physical Machines, Motherboards, NICs, or RAID Controllers

In the event of a hardware failure requiring replacement of an entire PM, everRun Enterprise supports PM replacement while VMs continue to run on the other PM thereby avoiding any guest downtime. The procedure for doing this is documented in ["Replacing Physical Machines, Motherboards, NICs, or RAID Controllers" on page 182](#). After completing this procedure, perform the following steps to avoid any potential subsequent problems.

Notes:



1. The physical machine **Replace** operation is used for a variety of reasons. After implementing the procedure described below, contact Stratus Support for guidance if there is any issue that does not resolve automatically, including any questionable indicators on the everRun Availability Console. For contact information see the **everRun Enterprise Downloads and Support** page at <http://www.stratus.com/go/support/everrun>.
2. Verify that the guest operating systems are running, then contact Stratus with all relevant status and questions, before attempting resolution on your own.

1. Open a terminal or ssh to the host OS on the PM that was **not** replaced. Log on as root and execute the following command:

```
pkill -x spine || echo "spine restart failed"
```

2. The preferred PM setting can get corrupted after replacing a PM. It is not possible to tell which VMs may be impacted by this, so you must perform steps **a.** and **b.** below for **each** VM.
 - a. In the everRun Availability Console, change the preferred PM setting to the node that the VM is **not** currently on. See ["Load Balancing" on page 113](#) and ["Selecting a Preferred PM for a Virtual Machine" on page 174](#) for information about how to do this.
 - b. Then change the preferred PM of each VM back to the desired setting.
3. If after performing the previous step an alert appears saying that a rebalance must be performed, on the everRun Availability Console masthead click **Rebalance** .

Error Message May Appear When Importing a VM to a HA-only Licensed System

When importing a VM to a system that is licensed to run only high availability (HA) VMs, deselect the option to **Auto start Virtual Machine after import**. Once the import completes, follow the instructions in ["Changing the Protection Level for a Virtual Machine \(HA or FT\)" on page 174](#) to reconfigure the VM's protection level to HA. Then click **Start** to start the VM.

If this option remains selected, the import will proceed but a message displays, stating that an error occurred. Contrary to the error message, the import operation actually succeeded. Follow the instructions in the previous paragraph to reconfigure and start the VM.

Using Internet Explorer Can't Change Volume Name During VM Import or Restore

When using the everRun Availability Console with Internet Explorer version 10.0.9200.16844, you cannot specify a new volume name while performing a VM import or restore operation.

Unsupported Network Adapter Card and Chip

Due to the issue outlined at <http://www-947.ibm.com/support/entry/portal/docdisplay?Indocid=migr-5093183>, everRun does not support the following network adapter card and chip:

- The Broadcom NetXtreme II Dual Port 10GBase-T Network Adapter, IBM part number 49Y7910
- Any other NICs which use the Broadcom BCM57712 Ethernet hardware chip.

Do Not Use the `ifdown` Command

Do not issue the `ifdown` command from an everRun Enterprise physical machine's host OS to temporarily bring down a VM's business (ibizx) network connection. 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.

New Features, Enhancements, and Bug Fixes

Major new features, enhancements, and bug fixes are listed below under the release in which they became available.

New in everRun Enterprise Release 7.1.0.0

- Official support of HA guests (was a technology preview in everRun Enterprise 7.0.x)
- Official support of split-site operation (was a technology preview in everRun Enterprise 7.0.x)
- Official support of quorum servers (was a technology preview in everRun Enterprise 7.0.x). Quorum server is an optional Windows-based service running on non-everRun Enterprise systems which increase overall system availability and protect against partition in time and split-brain scenarios.
- Support of 12 VMs (4 FT VMs maximum)
- Support for replacing a NIC or an everRun Enterprise physical machine without any guest VM

downtime

- Support for guests running the SuSE operating system
- Enhancement to present virtual SandyBridge family processor to guest OSs
- Improved presentation of virtual sockets and virtual CPUs to the guest to aid with (virtual) socket based software licenses
- Other serviceability, robustness, and usability improvements

Getting Help

If you have a technical question about everRun Enterprise software, you can find the latest documentation at <http://www.stratus.com/go/support/everrun>.

If you are unable to resolve your questions with the online documentation, and the system is covered by a service agreement, please contact everRun Enterprise Customer Support or your authorized Stratus service representative. For information, see the **everRun Enterprise Downloads and Support** page at <http://www.stratus.com/go/support/everrun>.

11

Chapter 11: everRun Enterprise Command Line Interface Reference

You can use the everRun Enterprise command-line interface to control the system from a remote console.

The following topics describe how to administer and use the command-line interface:

- ["AVCLI Command Overview" on page 194](#)
- ["AVCLI Command Descriptions" on page 206](#)

AVCLI Command Overview

You can use the everRun Enterprise command-line interface (AVCLI) to control the system from a remote console.

The following topics explain how to install the AVCLI client:

- ["Prerequisites" on page 195](#)
- ["Installing the Linux Client" on page 195](#)
- ["Installing the Windows Client" on page 196](#)

The following topics explain how to use the AVCLI command interface:

- ["Using AVCLI" on page 197](#)
- ["Executing a Command" on page 197](#)
- ["Using AVCLI Help" on page 198](#)

The following topics are helpful for programmers using the AVCLI command interface:

- ["AVCLI Error Status" on page 199](#)
- ["XML Encapsulated Errors" on page 199](#)
- ["Error Checking" on page 200](#)
- ["Asynchronous Command Delay" on page 200](#)
- ["Output Formatting" on page 201](#)
- ["AVCLI Exceptions" on page 205](#)

Related Topics

["AVCLI Command Descriptions" on page 206](#)

Prerequisites

Before you use AVCLI, these prerequisites apply:

- Verify that the client computer has Java Runtime Environment (JRE), version 1.6, update 14 or later, installed by typing:

```
java -version
```

If the client computer already has this version of JRE installed, the output looks similar to this:

```
java version "1.6.0_16" Java(TM) SE Runtime Environment
(build 1.6.0_16-b01) Java HotSpot(TM) Server VM (build
14.2-b01, mixed mode)
```

If the output shows that the client computer has an older version of JRE installed, download the correct version from <http://www.java.com/en/download/manual.jsp>.

- You need a valid username and password. The default username/password is `admin/admin`. AVCLI scripts embed the username/password, so use access control lists (ACLs) to safeguard your new credentials. AVCLI commands are encrypted using SSL.

Installing the Linux Client

To download the AVCLI client for Linux:

1. Download the Linux client:
 - a. Go to the **everRun Enterprise Downloads and Support** page at <http://www.stratus.com/go/support/everrun>.

- b. In the left-hand column, click **Drivers and Tools**.
 - c. Under **everRun Enterprise Command Line Interface (AVCLI)**, click **Download the RHEL 6 (64-bit) AVCLI Client**. Save the file.
2. Log on as the root user.
3. Add the directory `/usr/bin`, if it does not exist.
4. Install the client by typing:

```
rpm -i avcli*.rpm
```

Your Linux system can contain only one copy of AVCLI at one time. If another version is already installed, you receive an error message similar to the following:

```
file /usr/bin/avcli.bat from install of avcli-2.1.1-0 con-
flicts with file from package avcli-1.0-0 file /us-
r/lib/ImportExportLibs.jar from install of avcli-2.1.1-0
conflicts with file from package avcli-1.0-0
```

If you receive the preceding message, remove the previous version of AVCLI by typing:

```
rpm -e avcli-1.0-0
```

Then repeat step 4.

Installing the Windows Client

To download the AVCLI client for Windows:

1. Download the Windows client:
 - a. Go to the **everRun Enterprise Downloads and Support** page at <http://www.stratus.com/go/support/everrun>.
 - b. In the left-hand column, click **Drivers and Tools**.
 - c. Under **everRun Enterprise Command Line Interface (AVCLI)**, click **Download the Windows AVCLI Client**. Save the file.
2. Double-click `avcli.msi`. Follow the onscreen instructions.
3. Click **Run**. When prompted, accept the EULA.

4. If prompted to remove a previous version of AVCLI, click `Start > All Programs > everRun > Uninstall AVCLI`. Then repeat steps 1-3.

Using AVCLI

To use AVCLI:

- On Windows, click `Start Menu > All Programs > everRun > Command Prompt`.
- On Linux, type the `avcli` command, followed by one or more commands. For example:

```
# avcli -H localhost -u admin -p admin vm-info
```



Note: In the preceding example, typing the `-H`, `-u`, and `-p` options automatically saves the host-name, username, and password, respectively, so that subsequent commands do not require them. You can also create a shortcut to avoid the need to prefix all commands with the host-name, username, and password, as described in ["Executing a Command" on page 197](#).

From the command line, use the `help` command to list all AVCLI commands or to display information for a specific command. See ["Using AVCLI Help" on page 198](#).

Executing a Command

Commands must include the DNS name or IPv4 address of the everRun Enterprise system. If you specify incorrect syntax, a message displays the correct syntax.

Create a shortcut to avoid the need to prefix all commands with the hostname, username, and password.

To create a shortcut:

Windows

The `avcli` command executes the batch file `avcli.bat` in `%Program Files%\everRun`. You can add login credentials to this file:

1. Open `avcli.bat` with a text editor.
2. Search for this string:

```
-jar "%AVCLI_HOME%\avcli.jar"
```

3. Append login info. For example:

```
-jar "%AVCLI_HOME%\avcli.jar" -u admin -p admin -H everrun
```

If you manage several everRun systems with the same username and password, specify the domain names of the individual systems in the command line:

```
$ avcli -H everrun1 node-info node0
```

or

```
$ avcli -H everrun2 node-info node0
```

Linux

Create an alias in your login `.cshrc` file. For example:

```
alias avcli='/usr/bin/avcli -u admin -p admin -H everrun'
```

In the example, `avcli` is the alias name, `admin/admin` is the username/password, and `everrun` is the everRun system's domain name. You can then use this alias to log on and specify commands. For example, you could specify `unit-info` as follows:

```
$ avcli unit-info
```

Using AVCLI Help

This topic describes how to use AVCLI help.

Listing All Commands

To list all available AVCLI commands, type:

```
$ avcli help
```

The output follows:

```
[root@node0 zoneinfo]# avcli help
Usage: avcli [OPTION]... [COMMAND]
-u, --username username to login with
-p, --password password to login with
-H, --hostname hostname to login to
--log log file to capture debug information in
-x, --xml format output in XML
-V, --version display the version and exit
```

```
-h, --help display this message and exit  
  
.  
  
.  
  
.
```

If you type a command that AVCLI does not recognize, AVCLI displays the preceding output.

Displaying Help for a Specific Command

To display help for a specific command, type:

```
$ avcli help command_name
```

For example, if you type:

```
$ avcli help vm-create
```

The output is:

```
Usage: avcli vm-create[--interfaces] [--shared-storage]  
  
Create a new VM.  
  
.  
  
.  
  
.
```

If you type a valid command with an invalid argument, AVCLI displays the same output as if you had specified help for the command.

AVCLI Error Status

AVCLI does not follow the Linux convention of returning 0 on successful execution and 1 for an error.

XML Encapsulated Errors

To display all errors as encapsulated XML suitable for processing with an XML parser, specify `-x` on the command line.

The following example displays errors associated with a bad username/password:

```
$ avcli -x -H eagles -u admin -p foo node-info
```

The following example displays errors associated with a bad host address for the everRun Enterprise system:

```
$ avcli -x -H foo -u admin -p foo node-info

foo
```

The following example attempts an operation using a nonexistent VM:

```
$ avcli -H eagles -x vm-delete eagles23

Cannot find a resource that matches the identifier eagles23.
```

Error Checking

To cleanly catch all errors while developing scripts, always specify output in XML format. This returns an error with any reply that does not return valid XML or any XML document with an error attribute.

The following example is from a PERL subroutine `_cli` that provides a shell for executing AVCLI commands. The code that checks for errors does a simple pattern match on `$stdout`.

```
my $error = 0

$error = 1 unless ($stdout =~ /xml version/);

$error = 1 if ($stdout =~ /\//);
```

If no error occurs, `$stdout` is cast into a PERL hash using the standard PERL XML: :Simple Library. Otherwise, this error appears:

```
unless ($error) {

    my $xs = XML::Simple->new();

    $stdout_hash = $xs->XMLin($stdout,forceArray=>0);

    return 0;

}

return 1;
```

Asynchronous Command Delay

Commands that invoke an action on the everRun Enterprise system are called *asynchronous* because the command completes before the action completes. This allows for complex scripting.

If you want a command to complete inline before proceeding to the next command, create a simple script and use the `--wait` option. For example:

```
$ cli -x -H eagles node-workon --wait node0
```

In this example, `cli` does not complete until VMs and the management port are failed over from `node0` to `node1`, and `node0` is in maintenance mode. Without the `--wait` option, the command completes when acknowledged but before the resources are migrated.

Output Formatting

AVCLI can create user-friendly command output and program-friendly XML output.

User-Friendly Command Output

AVCLI output is formatted for easy readability. For example:

```
$ avance -u admin -p admin -H avance -x node-info
node:
-> name : node0
-> id : host:o14
-> state: running
-> sub-state : nil
-> standing-state : maintenance
-> mode : maintenance
-> primary : false
-> manufacturer : Dell
-> model : Dell PowerEdge 2950
-> maintenance-allowed : true
-> maintenance-guest-shutdown : false
-> cpus : 8
-> memory : 4,288,675,840
virtual machines:
```

```
node:
-> name : node1
-> id : host:o406
-> state : running
-> sub-state : nil
-> standing-state : warning
-> mode : normal
-> primary : true
-> manufacturer : Dell
-> model : Dell PowerEdge 2950
-> maintenance-allowed : true
-> maintenance-guest-shutdown : true
-> cpus : 8
-> memory : 4,288,675,840

virtual machines:
virtual machine:
-> name : eagles1
-> id : vm:o1836
```



Note: The output format of these commands may vary between releases.

Program-Friendly XML Output

You can create program-friendly XML output by using the `-x` or `--xml` global option. For example:

```
$ avcli -u admin -p admin -H localhost -x node-info
<?xml version="1.0" encoding="utf-8" standalone="no"?>
<avance>
<node>
```

```
<name>node1</name>
<id>host:o55</id>
<state>running</state>
<sub-state/>
<standing-state>normal</standing-state>
<mode>normal</mode>
<primary>>false</primary>
<manufacturer>Intel Corporation</manufacturer>
<model>S5520UR</model>
<maintenance-allowed>>true</maintenance-allowed>
<maintenance-guest-shutdown>>false</maintenance-guest-shutdown>
<cpus>2</cpus>
<memory>25706889216</memory>
<virtual-machines/>
</node>
<node>
<name>node0</name>
<id>host:o23</id>
<state>running</state>
<sub-state/>
<standing-state>normal</standing-state>
<mode>normal</mode>
<primary>>true</primary>
<manufacturer>Intel Corporation</manufacturer>
<model>S5520UR</model>
```

```

<maintenance-allowed>true</maintenance-allowed>
<maintenance-guest-shutdown>>false</maintenance-guest-shutdown>
<cpus>2</cpus>
<memory>25706889216</memory>
<virtual-machines>
<virtual-machine>
<name>MyVM</name>
<id>vm:o6417</id>
</virtual-machine>
</virtual-machines>
</node>
</avance>

```



Note: The schema definition is maintained between releases.

If you do **not** specify `-X` or `--XML` and the command returns an error, a verbose message appears. For example:

```

$ cli -H eagles vm-delete eagles23

%Error: Cannot find a resource that matches the identifier
eagles23. com.avance.yak.cli.exceptions.CommandLineException:
Cannot find a resource that matches the identifier eagles23.
at com.avance.yak.cli.ResourceDisambiguateServiceProvider.throwNonExistentResource
(ResourceDisambiguateServiceProvider.java:56)
at
com.avance.yak.cli.ResourceDisambiguateServiceProvider.getResourceId(ResourceDis-
ambiguateServiceProvider.java:81)

```

```
at
com.avance.yak.cli.Command.findResourceId(Command.java:80)
at com.avance.yak.cli.Com-
mandWithUnparsedAmbiguousResourcesInvokeEach.execute(Com-
mandWithUnparsedAmbiguousResourcesInvokeEach.java:65)
at
com.avance.yak.cli.Command.execute(Command.java:194)
at
com.avance.yak.cli.CommandLine.execute(CommandLine.java:649)
at
```

AVCLI Exceptions

If you do not specify `-X` or `--XML` and the command returns an error, a verbose message appears. For example:

```
$ cli -H eagles vm-delete eagles23

%Error: Cannot find a resource that matches the identifier
eagles23. com.avance.yak.cli.exceptions.CommandLineException:
Cannot find a resource that matches the identifier eagles23.
at com.avance.yak.cli.Re-
sourceDisambiguateServiceProvider.throwNonExistentResource
(ResourceDisambiguateServiceProvider.java:56)
at
com.avance.yak.cli.Re-
sourceDisambiguateServiceProvider.getResourceId(ResourceDis-
ambiguateServiceProvider.java:81)
at
com.avance.yak.cli.Command.findResourceId(Command.java:80)
```

```
at com.avance.yak.cli.Com-  
mandWithUnparsedAmbiguousResourcesInvokeEach.execute (Com-  
mandWithUnparsedAmbiguousResourcesInvokeEach.java:65)  
  
at  
  
com.avance.yak.cli.Command.execute (Command.java:194)  
  
at  
  
com.avance.yak.cli.CommandLine.execute (CommandLine.java:649)  
  
at  
  
com.avance.yak.cli.Program.main (Program.java:94)
```

AVCLI Command Descriptions

Click each heading to see the full list of AVCLI commands in that group.



Note: The Examples section for each command assumes that you have set up a command shortcut as described in ["Executing a Command" on page 197](#).

Help

["help" on page 233](#)

Basic System Information

["audit-export" on page 213](#)

["audit-info" on page 214](#)

["unit-configure" on page 288](#)

["unit-eula-accept" on page 289](#)

["unit-eula-reset" on page 290](#)

["unit-info" on page 291](#)

["unit-shutdown" on page 292](#)

["unit-shutdown-cancel" on page 293](#)

["unit-shutdown-state" on page 294](#)

["unit-synced" on page 295](#)

System Configuration

["callhome-disable" on page 215](#)

["callhome-enable" on page 216](#)

["callhome-info" on page 217](#)

["datetime-config" on page 218](#)

["dialin-disable" on page 226](#)

["dialin-enable" on page 227](#)

["dialin-info" on page 228](#)

["ealert-config" on page 229](#)

["ealert-disable" on page 230](#)

["ealert-enable" on page 231](#)

["ealert-info" on page 232](#)

["license-info" on page 237](#)

["license-install" on page 238](#)

["ntp-config" on page 271](#)

["ntp-disable" on page 272](#)

["proxy-config" on page 277](#)

["proxy-disable" on page 278](#)

["proxy-enable" on page 279](#)

["proxy-info" on page 280](#)

["snmp-config" on page 281](#)

["snmp-disable" on page 282](#)

["snmp-info" on page 283](#)

["timezone-config" on page 286](#)

["timezone-info" on page 287](#)

System User Management

["local-group-add" on page 239](#)

["local-group-delete" on page 240](#)

["local-group-edit" on page 241](#)

["local-group-info" on page 242](#)

["local-user-add" on page 243](#)

["local-user-delete" on page 245](#)

["local-user-edit" on page 246](#)

["local-user-info" on page 248](#)

["owner-config" on page 275](#)

["owner-info" on page 276](#)

Managing Physical Machines

["node-add" on page 259](#)

["node-cancel" on page 260](#)

["node-delete" on page 261](#)

["node-info" on page 262](#)

["node-poweroff" on page 263](#)

["node-poweron" on page 264](#)

["node-reboot" on page 265](#)

["node-recover" on page 266](#)

["node-shutdown" on page 267](#)

["node-upgrade" on page 268](#)

["node-workoff" on page 269](#)

["node-workon" on page 270](#)

Managing Alerts

["alert-delete" on page 211](#)

["alert-info" on page 212](#)

Diagnostic Files

["diagnostic-create" on page 221](#)

["diagnostic-delete" on page 222](#)

["diagnostic-extract" on page 223](#)

["diagnostic-fetch" on page 224](#)

["diagnostic-info" on page 225](#)

["kit-delete" on page 234](#)

["kit-info" on page 235](#)

["kit-upload" on page 236](#)

Network/Storage Information

["network-change-mtu" on page 255](#)

["network-change-role" on page 256](#)

["network-info" on page 257](#)

["storage-group-info" on page 284](#)

["storage-info" on page 285](#)

["volume-info" on page 314](#)

Creating Virtual CD/DVDs

["media-create" on page 249](#)

["media-delete" on page 250](#)

["media-eject" on page 251](#)

["media-import" on page 252](#)

["media-info" on page 254](#)

Managing Virtual Machines

["ova-info" on page 273](#)

["ovf-info" on page 274](#)

["vm-boot-attributes" on page 296](#)

["vm-cd-boot" on page 297](#)

["vm-create" on page 298](#)

["vm-delete" on page 301](#)

["vm-import" on page 302](#)

["vm-info" on page 305](#)

["vm-migrate" on page 306](#)

["vm-poweroff" on page 307](#)

["vm-poweron" on page 308](#)

["vm-reprovision" on page 309](#)

["vm-shutdown" on page 312](#)

["vm-unlock" on page 313](#)

Related Topics

["AVCLI Command Overview" on page 194](#)

alert-delete

Usage

```
avcli alert-delete [alerts... | purge]
```

Description

The `alert-delete` command deletes specific alerts or optionally, all alerts.

Options

<i>alerts</i>	One or more alerts to be deleted.
<code>purge</code>	Delete all alerts.

Examples

```
$ avcli alert-delete alert:o10
```

```
$ avcli alert-delete alert:o10 alert:o11
```

```
$ avcli alert-delete purge
```

alert-info**Usage**

```
avcli alert-info [alerts...]
```

Description

The `alert-info` command displays information about all alerts or only those specified.

Options

<i>alerts</i>	The alert information to be displayed.
---------------	----------------------------------------

audit-export

Usage

```
avcli audit-export
```

Description

The `audit-export` command exports all of the audit logs.

audit-info**Usage**

```
avcli audit-info [number-of-audit-logs]
```

Description

The `audit-info` command displays either the last 50 audit logs or the specified number of audit logs.

Options

<i>number-of-audit-logs</i>	The number of audit logs to display. The default value is 50.
-----------------------------	---------------------------------------------------------------

Examples

```
$ avcli audit-info
```

```
$ avcli audit-info 25
```

callhome-disable

Usage

```
avcli callhome-disable
```

Description

The `callhome-disable` command disables call home.

callhome-enable

Usage

```
avcli callhome-enable
```

Description

The `callhome-enable` command enables call home.

callhome-info

Usage

```
avcli callhome-info
```

Description

The `callhome-info` command displays information about call home.

datetime-config**Usage**

```
avcli datetime-config date time [timezone]
```

Description

The `datetime-config` command sets the date, time, and time zone on everRun systems.

Options

<i>date</i>	The date, formatted as <i>YYYY-MM-DD</i> .
<i>time</i>	The time, formatted as <i>HH:MM:SS</i> , in a 24-hour format.
<i>timezone</i>	The time zone. By default, it is the currently configured time zone.

You can specify the following values for *timezone*.

Africa/Cairo	Africa/Casablanca	Africa/Harare
Africa/Lagos	Africa/Monrovia	Africa/Nairobi
Africa/Windhoek	America/Adak	America/Anchorage
America/Asuncion	America/Bogota	America/Buenos_Aires
America/Caracas	America/Chicago	America/Chihuahua
America/Cuiaba	America/Denver	America/Godthab
America/Goose_Bay	America/Grand_Turk	America/Guyana
America/Halifax	America/Havana	America/Indianapolis
America/Los_Angeles	America/Managua	America/Manaus
America/Mexico_City	America/Miquelon	America/Montevideo

America/New_York	America/Noronha	America/Phoenix
America/Regina	America/Santiago	America/Sao_Paulo
America/St_Johns	America/Tijuana	America/Winnipeg
Asia/Amman	Asia/Baghdad	Asia/Baku
Asia/Bangkok	Asia/Beijing	Asia/Beirut
Asia/Bishkek	Asia/Calcutta	Asia/Colombo
Asia/Damascus	Asia/Dhaka	Asia/Gaza
Asia/Hong_Kong	Asia/Irkutsk	Asia/Jerusalem
Asia/Kabul	Asia/Kamchatka	Asia/Karachi
Asia/Katmandu	Asia/Krasnoyarsk	Asia/Magadan
Asia/Novosibirsk	Asia/Rangoon	Asia/Riyadh
Asia/Seoul	Asia/Singapore	Asia/Taipei
Asia/Tashkent	Asia/Tbilisi	Asia/Tehran
Asia/Tokyo	Asia/Vladivostok	Asia/Yakutsk
Asia/Yekaterinburg	Asia/Yerevan	Atlantic/Azores
Atlantic/Cape_Verde	Atlantic/Stanley	Australia/Adelaide
Australia/Brisbane	Australia/Darwin	Australia/Hobart
Australia/Lord_Howe	Australia/Melbourne	Australia/Perth
Australia/Sydney	Etc/GMT	Etc/GMT+1
Etc/GMT+10	Etc/GMT+11	Etc/GMT+12

Etc/GMT+2	Etc/GMT+3	Etc/GMT+4
Etc/GMT+5	Etc/GMT+6	Etc/GMT+7
Etc/GMT+8	Etc/GMT+9	Etc/GMT-1
Etc/GMT-10	Etc/GMT-11	Etc/GMT-12
Etc/GMT-13	Etc/GMT-14	Etc/GMT-2
Etc/GMT-3	Etc/GMT-4	Etc/GMT-5
Etc/GMT-6	Etc/GMT-7	Etc/GMT-8
Etc/GMT-9	Europe/Athens	Europe/Belgrade
Europe/Berlin	Europe/Helsinki	Europe/Istanbul
Europe/Kaliningrad	Europe/London	Europe/Minsk
Europe/Moscow	Europe/Paris	Europe/Samara
Europe/Sarajevo	Japan	Pacific/Auckland
Pacific/Chatham	Pacific/Easter	Pacific/Fiji
Pacific/Guam	Pacific/Marquesas	Pacific/Norfolk
Pacific/Tongatapu		

Examples

```
$ avcli datetime-config 2010-12-31 6:03:10
```

```
$ avcli datetime-config 2010-12-31 20:09:22 America/New_York
```

diagnostic-create

Usage

```
avcli diagnostic-create [minimal | medium | stats | full]
```

Description

The `diagnostic-create` command creates a new diagnostic of the specified type.

Options

<code>minimal</code>	The smallest diagnostic (approximately 2 to 10 MB).
<code>medium</code>	A medium diagnostic (approximately 10 MB).
<code>stats</code>	A medium diagnostic that includes statistics.
<code>full</code>	A large diagnostic (approximately 60 MB).

diagnostic-delete**Usage**

```
avcli diagnostic-delete diagnostics...
```

Description

The `diagnostic-delete` command deletes the specified diagnostic files.

Options

<i>diagnostics</i>	One or more diagnostic files to be deleted.
--------------------	---------------------------------------------

diagnostic-extract

Usage

```
avcli diagnostic-extract diagnostics.zip...
```

Description

The `diagnostic-extract` command extracts the specified diagnostic files.

Options

<i>diagnostics</i>	One or more diagnostic files to be extracted.
--------------------	-----------------------------------------------

diagnostic-fetch**Usage**

```
avcli diagnostic-fetch [--file name] diagnostics...
```

Description

The `diagnostic-fetch` command downloads the specified diagnostics to the current directory. If the diagnostic's status is `busy`, `diagnostic-fetch` waits for the diagnostic to complete and then downloads it. The default diagnostic file name is `diagnostic-type-name_YYYYMMDD_HHMMSS.zip`:

- *type*: The type of diagnostic: minimal, medium, stats, full.
- *name*: The name of the everRun system, as displayed by `unit-info`.
- *YYYY*: The year the diagnostic was created.
- *MM*: The month the diagnostic was created.
- *DD*: The day of the month the diagnostic was created.
- *HH*: The hour the diagnostic was created.
- *MM*: The minute the diagnostic was created.
- *SS*: The second the diagnostic was created.

Options

<i>diagnostics</i>	One or more diagnostic files to be downloaded.
<code>--file <i>name</i></code>	The name of the file written to the current directory. This option is valid if only one diagnostic is downloaded.
<code>--extract</code>	Extract the downloaded diagnostic files.

Examples

```
$ avcli diagnostic-fetch buggrab:o10
```

```
$ avcli diagnostic-fetch --file buggrab.zip buggrab:o10
```

```
$ avcli diagnostic-fetch buggrab:o10 buggrab:o11 buggrab:o12
```

diagnostic-info

Usage

```
avcli diagnostic-info diagnostics...
```

Description

The `diagnostic-info` command displays information about all diagnostics or only about those specified.

Options

<i>diagnostics</i>	One or more diagnostic files about which to display information.
--------------------	------------------------------------------------------------------

dialin-disable

Usage

```
avcli dialin-disable
```

Description

The `dialin-disable` command disables dial-in.

dialin-enable

Usage

```
avcli dialin-enable
```

Description

The `dialin-enable` command enables dial-in.

dialin-info

Usage

```
avcli dialin-info
```

Description

The `dialin-info` command displays information about dial-in configuration.

ealert-config

Usage

```
avcli ealert-config [--ssl] [--username name] [--password pass-
word] --host recipients...
```

Description

The `ealert-config` command configures e-Alert support in everRun systems. If you do not provide a user name, the command assumes that authentication is not required to access the SMTP server. If you provide a user name but not a password, the command prompts you for the password.

Options

<code>--ssl</code>	Use SSL when communicating with the SMTP server.
<code>--username <i>name</i></code>	The name to authenticate with against the specified SMTP host.
<code>--password <i>password</i></code>	The password to authenticate with against the specified SMTP host.
<code>--host <i>recipients</i></code>	The DNS or IP address of the SMTP server.

Examples

```
$ avcli ealert-config --host mail.my-domain.com admin@my-
domain.com
```

```
$ avcli ealert-config --host mail.my-domain.com admin@my-
domain.com bob@my-domain.com
```

```
$ avcli ealert-config --host mail.my-domain.com --username
admin --password secret --ssl bob@my-domain.com
```

```
$ avcli ealert-config --host mail.my-domain.com --username
admin --ssl bob@my-domain.com
```

ealert-disable

Usage

```
avcli ealert-disable
```

Description

The `ealert-disable` command disables e-Alert.

ealert-enable

Usage

```
avcli ealert-enable
```

Description

The `ealert-enable` command enables e-Alert.

ealert-info

Usage

```
avcli ealert-info
```

Description

The `ealert-info` command displays information about the e-Alert configuration.

help

Usage

```
avcli help [command] [-all]
```

Description

The `help` command provides help about a specific command or lists all AVCLI commands.

Options

<code>-all</code>	Display detailed information for all commands.
-------------------	------------------------------------------------

Examples

To display general command usage and a list of all commands for which `help` provides information:

```
$ avcli help
```

To display information about a specific command (`storage-info`, in this example):

```
$ avcli help storage-info
```

To display detailed information about all commands for which `help` provides information:

```
$ avcli help -all
```

kit-delete**Usage**

```
avcli kit-delete kit...
```

Description

The `kit-delete` command deletes the specified kits.

Options

<i>kit</i>	One or more upgrade kits to be deleted.
------------	-----------------------------------------

kit-info

Usage

```
avcli kit-info [kit...]
```

Description

The `kit-info` command displays information about all kits (the default) or only the specified kits.

Options

<i>kit</i>	One or more upgrade kits about which to display information.
------------	--------------------------------------------------------------

kit-upload

Usage

```
avcli kit-upload kit...
```

Description

The `kit-upload` command uploads the specified kit files.

Options

<i>ki-</i> <i>t</i>	One or more upgrade kits to be uploaded.
------------------------	------------------------------------------

Examples

```
$ avcli kit-upload /var/kits/kit-avance.tar.bz2
```

license-info

Usage

```
avcli license-info
```

Description

The `license-info` command displays information about the license.

license-install

Usage

```
avcli license-install license-file
```

Description

The `license-install` command installs the specified license file.

Options

<i>license-file</i>	The file containing the license-key definitions.
---------------------	--------------------------------------------------

Examples

```
$ avcli license-install avance.key
```

local-group-add

Usage

```
avcli local-group-add --name name --permissions permission-type
```

Description

The `local-group-add` command adds a new local user group.

Options

<code>--name <i>name</i></code>	Local group name.
<code>--permissions <i>permission-type</i></code>	Local group permissions, in the form of a comma-separated list.

Examples

```
$ avcli local-group-add --name unprivileged_users --permissions ADD_USER
```

local-group-delete

Usage

```
avcli local-group-delete groups...
```

Description

The `local-group-delete` command deletes the specified local user groups. You cannot delete default groups (`admin`, `platform_admin`, `read_only`).

Options

<i>groups</i>	Local user groups.
---------------	--------------------

Examples

```
$ avcli local-group-delete unprivileged_users
```

local-group-edit

Usage

```
avcli local-group-edit [--name] [--permissions] group-name-or-sid
```

Description

The `local-group-edit` command edits an existing local user group. You cannot edit default groups (`admin`, `platform_admin`, `read_only`).

Options

<code>--name</code> <i>name</i>	New local group name.
<code>--permissions</code> <i>permission-type</i>	Local group permissions, in the form of a comma-separated list.
<i>group-name-or-sid</i>	The name or security ID.

Examples

```
$ avcli local-group-edit --name privileged_users --permissions  
ADD_USER unprivileged_users
```

local-group-info**Usage**

```
avcli local-group-info [groups...]
```

Description

The `local-group-info` command displays information about all local user groups, or only those specified.

Options

<i>groups</i>	Local user groups.
---------------	--------------------

local-user-add

Usage

```
avcli local-user-add --username name --realname name --email
address [--password password] [--new-password password] [--
local-groups groups] [--permissions permission-types]
```

Description

The `local-user-add` command adds a new local user to the everRun system. If the user's password is not given, the user is automatically prompted for it. The user is prompted twice to verify that the password was entered correctly.

Options

<code>--username <i>name</i></code>	everRun local user name.
<code>--password <i>password</i></code>	Boolean flag that indicates whether the user should be prompted for a new password.
<code>--new-password <i>password</i></code>	Specify the password as a command-line option instead of being prompted in the same way as <code>--password</code> .
<code>--realname <i>name</i></code>	The user's real name.
<code>--email <i>address</i></code>	The user's email address.
<code>--local-groups <i>groups</i></code>	Local groups for the user to join, in the form of a comma-separated list.
<code>--permissions <i>permission-types</i></code>	Local user permissions, in the form of a comma-separated list.

Examples

```
$ avcli local-user-add --username bsmith --realname "Bob  
Smith" --email bsmith@example.com --password secret --local-  
groups admin
```

```
$ avcli local-user-add --username bsmith --realname "Bob  
Smith" --email bsmith@example.com --local-groups users1,users2  
--permissions ADD_USER,UPDATE_USER
```

local-user-delete

Usage

```
avcli local-user-delete users...
```

Description

The `local-user-delete` command deletes the specified local users.

Options

<i>users</i>	One or more local users.
--------------	--------------------------

Examples

```
$ avcli local-user-delete afjord
```

```
$ avcli local-user-delete afjord bsmith tkirch
```

local-user-edit**Usage**

```
avcli local-user-edit user [--username name] [--realname
name] [--email address] [--password password] [--new-password
password] [--local-groups groups] [--permissions permission-
types] user-name-or-sid
```

Description

The `local-user-edit` command edits an existing user. If you do not give the `--password` option, the password is not changed. If you give the `--password` option, the command prompts the user twice to verify that the password was entered correctly.

Options

<code>--username <i>name</i></code>	The user name to assign.
<code>--password <i>password</i></code>	Boolean flag that indicates whether the user should be prompted for a new password.
<code>--new-password <i>password</i></code>	Specify the password as a command-line option instead of being prompted in the same way as <code>--password</code> .
<code>--realname <i>name</i></code>	The user's real name.
<code>--email <i>address</i></code>	The user's email address.
<code>--local-groups <i>groups</i></code>	Local groups for the user to join, in the form of a comma-separated list.
<code>--permissions <i>permission-types</i></code>	Local user permissions, in the form of a comma-separated list.
<code><i>group-name-or-sid</i></code>	The name or security ID.

Examples

```
$ avcli local-user-edit --email bsmith@example.net bsmith  
$ avcli local-user-edit --realname "Robert Smith" --email  
rsmith@example.com bsmith  
$ avcli local-user-edit --email bsmith@example.net --local-  
groups read_only --permissions ADD_USER,UPDATE_USER bsmith  
$ avcli local-user-edit --password bsmith  
$ avcli local-user-edit --new-password secret bsmith
```

local-user-info**Usage**

```
avcli local-user-info [user...]
```

Description

The `local-user-info` command displays information about all users (by default) or only about the specified users.

Options

<i>user</i>	One or more users about which to display information.
-------------	-------------------------------------------------------

media-create

Usage

```
avcli media-create [--storage-group storage] [--name name]
url...
```

Description

The `media-create` command loads an ISO image into an everRun system from the specified URL.

Options

<code>--storage-group <i>group</i></code>	The storage volume to carve from. If you do not specify this option, the storage group with the most free space is automatically selected.
<code>--name <i>name</i></code>	The name of the carved volume. If you do not specify this option, the name is determined from the URL.
<code><i>url</i></code>	The URL where the ISO file is located.
<code>--wait</code>	Wait for the ISO(s) to be created.

Examples

```
avcli media-create --storage-group Pool-0001 --name cd.iso
http://hostname/cd.iso

avcli media-create http://hostname/cd.iso

avcli media-create http://hostname/cd1.iso http://host-
name/cd2.iso
```

media-delete**Usage**

```
avcli media-delete media...
```

Description

The `media-delete` command deletes the specified media.

Options

<i>media</i>	The media to be deleted.
--------------	--------------------------

media-eject

Usage

```
avcli media-eject [--cdrom name] [vm...]
```

Description

The `media-eject` command ejects media from the specified virtual machines.

Options

<code>--cdrom <i>name</i></code>	The CD-ROM device to eject. This value is optional if the VM has only a single CD-ROM device.
<code><i>vm</i></code>	The name of the VM containing the media to be ejected.

media-import**Usage**

```
avcli media-import [--storage-group storage] [--name name] [--throttle] [--silent] file...
```

Description

The `media-import` command loads an ISO image into an everRun system from the specified file.

Options

<code>--storage-group <i>group</i></code>	The storage volume to carve from. If you do not specify this option, the shared storage with the most free space is automatically selected.
<code>--name <i>name</i></code>	The name of the carved volume. If you do not specify this option, the name is determined from the file. This option is valid only if one ISO is specified.
<code>--throttle</code>	Slow down the import/export operation. Valid values are: <ul style="list-style-type: none"> • <code>none</code>: Do not use throttling. This is the default value. • <code>low</code>: Slow down by about 25%. • <code>medium</code>: Slow down by about 50%. • <code>high</code>: Slow down by about 75%.
<code>--silent</code>	Suppress output.
<i>file</i>	The file(s) containing an ISO image.

Examples

```
avcli media-import --storage-group Pool-0001 --name cd.iso
cd.iso
```

```
avcli media-import cd.iso
```

```
avcli media-import cd1.iso cd2.iso
```

media-info

Usage

```
avcli media-info [media...]
```

Description

The `media-info` command displays information about all media, or optionally, only the specified media.

Options

<i>media</i>	The media about which to display information.
--------------	-----------------------------------------------

network-change-mtu

Usage

```
avcli network-change-mtu name size
```

Description

The `network-change-mtu` command changes the MTU size of the specified A-Link network on everRun systems.

Options

<i>name</i>	The name of the A-Link network
<i>size</i>	The MTU size. Valid values are 1500 - 9000.

Examples

```
$ avcli network-change-mtu priv0 4000
```

```
$ avcli network-change-mtu priv0 9000
```

network-change-role

Usage

```
avcli network-change-role networks... role
```

Description

The `network-change-role` command changes the role of the specified network to the specified role.

Options

<i>networks</i>	One or more networks whose role is to be changed.
<i>role</i>	The new role. Specify either <code>business</code> or <code>a-link</code> .

network-info

Usage

```
avcli network-info [networks...]
```

Description

The `network-info` command displays information about all shared networks, or optionally, about only the specified networks.

Options

<i>networks</i>	One or more networks.
-----------------	-----------------------

Output

The following example shows the settings for four networks, including the default MTU value of 1500 for A-Links.

```
avcli network-info
shared network:
  -> name           : sync_2003
  -> id             : sharednetwork:o2334
  -> fault-tolerant : ft
  -> role           : a-link
  -> bandwidth      : 10 Gb/s
  -> mtu            : 1500
shared network:
  -> name           : network0
  -> id             : sharednetwork:o64
  -> fault-tolerant : ft
  -> role           : business
  -> bandwidth      : 1 Gb/s
  -> mtu            : 1500
```

shared network:

```
-> name          : sync_2004
-> id            : sharednetwork:o2333
-> fault-tolerant : ft
-> role          : a-link
-> bandwidth     : 10 Gb/s
-> mtu           : 1500
```

shared network:

```
-> name          : priv0
-> id            : sharednetwork:o65
-> fault-tolerant : ft
-> role          : private
-> bandwidth     : 1 Gb/s
-> mtu           : 1500
```

node-add

Usage

```
avcli node-add [--wait]
```

Description

The `node-add` command adds a PM to an everRun system.

Options

<code>--wait</code> <code>-w</code>	Wait for the command to complete.
----------------------------------------	-----------------------------------

node-cancel**Usage**

```
avcli node-cancel pm
```

Description

The `node-cancel` command cancels a PM that is being imaged.

Options

<i>pm</i>	The PM to cancel.
-----------	-------------------

node-delete

Usage

```
avcli node-delete pm [--wait]
```

Description

The `node-delete` command deletes a PM.

Options

<i>pm</i>	The PM to be deleted. It must be in maintenance mode.
--wait -w	Wait for the command to complete.

node-info**Usage**

```
avcli node-info [pm...]
```

Description

The `node-info` command displays information about all PMs (by default) or only about the specified PMs.

Options

<i>pm</i>	The PM about which to display information.
-----------	--------------------------------------------

node-poweroff

Usage

```
avcli node-poweroff pm [--wait]
```

Description

The `node-poweroff` command powers off the specified PM.

Options

<i>pm</i>	The PM to be powered off.
--wait -w	Wait for the command to complete.

node-poweron**Usage**

```
avcli node-poweron pm [--wait]
```

Description

The `node-poweron` command powers on the specified PM.

Options

<i>pm</i>	The PM to be powered on.
--wait -w	Wait for the command to complete.

node-reboot

Usage

```
avcli node-reboot pm [--wait]
```

Description

The `node-reboot` command reboots the specified PM.

Options

<i>pm</i>	The PM to reboot.
--wait -w	Wait for the command to complete.

node-recover**Usage**

```
avcli node-recover [--wipe] pm [--wait]
```

Description

The `node-recover` command recovers the specified PM.

Options

<i>pm</i>	The PM to recover.
<code>--wipe</code>	Wipe the disks off the PM prior to recovery.
<code>--wait</code> <code>-w</code>	Wait for the command to complete.

node-shutdown

Usage

```
avcli node-shutdown pm [--force] [--wait]
```

Description

The `node-shutdown` command shuts down the specified PM.

Options

<i>pm</i>	The PM to shut down.
--force -f	Override the shutdown warning.
--wait -w	Wait for the command to complete.

node-upgrade

Usage

```
avcli node-upgrade --kit kit pm
```

Description

The `node-upgrade` command upgrades the PM with the specified kit.

Options

<i>pm</i>	The PM to upgrade.
<code>--kit <i>kit</i></code>	The kit to use for upgrading.

node-workoff

Usage

```
avcli node-workoff pm [--wait]
```

Description

The `node-workoff` command takes the specified PM out of maintenance mode.

Options

<i>pm</i>	The PM to remove from maintenance mode.
--wait -w	Wait for the command to complete.

node-workon**Usage**

```
avcli node-workon pm
```

Description

The `node-workon` command places the specified PM in maintenance mode.

Options

<i>pm</i>	The PM to place into maintenance mode.
-----------	----------------------------------------

ntp-config

Usage

```
avcli ntp-config servers...
```

Description

The `ntp-config` command enables and configures NTP support using the specified list of servers.

Options

<code>servers</code>	The list of servers to configure.
----------------------	-----------------------------------

Examples

```
$ avcli ntp-config 1.2.3.4
```

```
$ avcli ntp-config 1.2.3.4 2.4.6.8
```

ntp-disable

Usage

```
avcli ntp-disable
```

Description

The `ntp-disable` command disables NTP in your everRun system.

ova-info

Usage

```
avcli ova-info filename.ova...
```

Description

The `ova-info` command displays information about the specified OVA files.

Options

<code>filename.ova</code>	One or more OVA files.
---------------------------	------------------------

ovf-info**Usage**

```
avcli ovf-info filename.ovf...
```

Description

The `ovf-info` command displays information about the specified OVF files.

Options

<code><i>filename.ovf</i></code>	One or more OVF files.
----------------------------------	------------------------

owner-config

Usage

```
avcli owner-config [--email address] [--name name] [--phone  
number]
```

Description

The `owner-config` command configures the everRun system's owner information.

Options

<code>--email <i>address</i></code>	The owner's email address.
<code>--name <i>name</i></code>	The owner's name.
<code>--phone <i>number</i></code>	The owner's phone number.

Examples

```
$ avcli owner-config --email "Bob Smith" --email  
bsmith@example.org --phone 800-555-1234  
  
$ avcli owner-config --phone 800-555-1234
```

owner-info

Usage

```
avcli owner-info
```

Description

The `owner-info` command displays information about the owner of the everRun system.

proxy-config

Usage

```
avcli proxy-config --port name [--username name] [--password
password] host
```

Description

The `proxy-config` command configures the everRun system to use a proxy server. If you do not specify a user name, AVCLI assumes that authentication is not required to access the proxy server. If you specify a user name but not a password, you are prompted for the password.

Options

<code>--port <i>number</i></code>	The port number.
<code>--username <i>name</i></code>	The user's name.
<code>--password <i>password</i></code>	The user's password.
<code><i>host</i></code>	The host name.

Examples

```
$ avcli --port 8080 proxy.my-domain.com
```

```
$ avcli --port 8080 --username user --password secret
proxy.my-domain.com
```

```
$ avcli --port 8080 --username user proxy.my-domain.com
```

proxy-disable

Usage

```
avcli proxy-disable
```

Description

The `proxy-disable` command disables proxy.

proxy-enable

Usage

```
avcli proxy-enable
```

Description

The `proxy-enable` command enables proxy.

proxy-info

Usage

```
avcli proxy-info
```

Description

The `proxy-info` command displays information about the proxy configuration.

snmp-config

Usage

```
avcli snmp-config [--enable-requests] [--enable-traps] [--port
number] [--community name] [recipients...]
```

Description

The `snmp-config` command configures SNMP for use in an everRun system.

Options

<code>--enable-requests</code>	Enable SNMP requests. If you do not specify the option, requests are disabled.
<code>--enable-traps</code>	Enable SNMP traps. If you do not specify the option, traps are disabled.
<code>--community name</code>	The name of the SNMP community.
<code>--port number</code>	The port to use for SNMP. The default is 162.
<code>recipients</code>	The list of hosts to which to send traps; required only when traps are enabled.

Examples

The following example enables SNMP requests, and then traps and sends them to `localhost` and `snmp.my-domain.com`.

```
$ avcli snmp-config --enable-requests --enable-traps --com-
community public localhost snmp.my-domain.com
```

The following example disables SNMP requests, enables traps, and sends them to `localhost`.

```
$ avcli snmp-config --enable-traps --community public loc-
alhost
```

snmp-disable

Usage

```
avcli snmp-disable
```

Description

The `snmp-disable` command disables SNMP.

snmp-info

Usage

```
avcli snmp-info
```

Description

The `snmp-info` command displays information about SNMP configuration.

storage-group-info**Usage**

```
avcli storage-group-info [--disks] [--volumes] [storage-  
group...]
```

Description

The `storage-group-info` command displays information about all storage groups, or optionally, only about those specified.

Options

<code>--disks</code>	Show the logical disks belonging to a storage group.
<code>--volumes</code>	Show the volumes using a storage group.
<i>storage-group</i>	One or more storage groups about which to display information.

storage-info

Usage

```
avcli storage-info [--disks] [--volumes] [storage-group...]
```

Description

The `storage-info` command displays information about all storage groups, or optionally, only about those specified.

Options

<code>--disks</code>	Show the logical disks belonging to a storage group.
<code>--volumes</code>	Show the volumes using a storage group.
<i>storage-group</i>	One or more storage groups about which to display information.

timezone-config**Usage**

```
avcli timezone-config timezone
```

Description

The `timezone-config` command sets the time zone.

Options

<i>timezone</i>	The time zone.
-----------------	----------------

Examples

```
$ avcli timezone-config America/New_York
```

timezone-info

Usage

```
avcli timezone-info
```

Description

The `timezone-info` command displays the list of configurable time zones.

unit-configure

Usage

```
avcli unit-configure
```

Description

The `unit-configure` command configures the everRun system.

unit-eula-accept

Usage

```
avcli unit-eula-accept [--deny]
```

Description

The `unit-eula-accept` command accepts or denies the EULA.

Options

<code>--deny</code>	Deny acceptance of the EULA.
---------------------	------------------------------

unit-eula-reset

Usage

```
avcli unit-eula-reset
```

Description

The `unit-eula-reset` command resets the EULA acceptance state on an everRun system.

unit-info

Usage

```
avcli unit-info
```

Description

The `unit-info` command displays information about the specified everRun system.

unit-shutdown

Usage

```
avcli unit-shutdown
```

Description

The `unit-shutdown` command shuts down an everRun system.

unit-shutdown-cancel

Usage

```
avcli unit-shutdown-cancel
```

Description

The `unit-shutdown-cancel` command cancels a pending shutdown for an everRun system.

unit-shutdown-state

Usage

```
avcli unit-shutdown-state
```

Description

The `unit-shutdown-state` command returns the everRun system's shutdown state.

unit-synced

Usage

```
avcli unit-synced [--wait]
```

Description

The `unit-synced` command returns true if the everRun system is synchronized between all PMS; otherwise, it returns false.

Options

<code>--wait</code>	Wait for the command to complete.
<code>-w</code>	

vm-boot-attributes**Usage**

```
avcli vm-boot-attributes --priority priority --application-
start-time minutes [vm...]
```

Description

The `vm-boot-attributes` command sets the boot attributes for the specified VMs.

Options

<code>--priority <i>priority</i></code>	The boot priority; values are 1 to 1000.
<code>--application-start-time <i>minutes</i></code>	The estimated start time of the VM and application, in minutes. The minimum value is one minute.
<code><i>vm</i></code>	One or more VMs whose boot attributes are being set.

Examples

```
$ avcli vm-boot-attributes --priority 1 --application-start-
time 1 vm1
```

```
$ avcli vm-boot-attributes --priority 1 --application-start-
time 1 vm:o100
```

vm-cd-boot

Usage

```
avcli vm-cd-boot --iso iso [--wait] [vm...]
```

Description

The `vm-cd-boot` command starts the specified VMs and boots from the specified ISO image.

Options

<code>--iso <i>iso</i></code>	The ISO image to boot.
<code>--wait</code>	Wait for the VM to boot.
<code><i>vm</i></code>	One or more VMs that are being started.

Examples

```
$ avcli vm-cd-boot --iso MyISO vm1
```

```
$ avcli vm-cd-boot --iso MyISO vm:o100
```

```
$ avcli vm-cd-boot --iso MyISO --wait vm1
```

vm-create**Usage**

```
avcli vm-create --name name --cpu number --memory memory --
cdrom cd-name | --kickstart template [--interfaces networks]
[--storage-group group] --volumes volumes [--wait]
```

Description

The `vm-create` command creates a new VM.

Options

<code>--name <i>name</i></code>	The name of the VM to create.
<code>--cpu <i>number</i></code>	The number of virtual CPUs to assign to the VM.
<code>--memory <i>memory</i></code>	The amount of memory, in megabytes, to assign to the VM.
<code>--cdrom <i>cd-name</i></code>	The CD-ROM from which you initially boot the VM. You cannot specify this option with <code>--kickstart</code> .
<code>--kickstart <i>template</i></code>	The kickstart template to use when booting the VM. You cannot specify this option with <code>--cdrom</code> .
<code>--interfaces <i>networks</i></code>	The list of networks to attach to the VM. Specify a network only once. The attached network must not be private.
<code>--storage-group <i>group</i></code>	The storage group to use to carve VM volumes from. If you do not specify this value, the storage group with the most free space is automatically selected.
<code>--volumes <i>volumes</i></code>	The list of volumes to attach to this VM. A <i>volume</i> is made up of three components, separated by commas: <ul style="list-style-type: none"> • The size of the volume; required. • The storage-group name of the ID from which to carve storage.

	<ul style="list-style-type: none"> • The volume name. <p>By default, the volume size is specified in megabytes, but you can use standard qualifiers such as KB, MB, GB, and TB.</p>
<pre>--wait -w</pre>	<p>Wait for the command to complete.</p>

Examples

Create a VM named `vm001` with a CPU, 512 MB of memory, a 1,024 MB volume, and that is attached to `network0`.

```
$ avcli vm-create --name vm001 --cpu 1 --memory 512 --
cdrom linux.iso --interfaces network0 \ --volumes 1024
```

Create a VM named `vm001` with a CPU, 512 MB of memory, a 1,024 MB volume, and that is attached to `network0`. Then allocate storage from `Pool-0001` for the volume.

```
$ avcli vm-create --name vm001 --cpu 1 --memory 512 --
cdrom linux.iso --interfaces network0 \ --volumes 1024 --
storage-group Pool-0001
```

Create a VM named `vm001` with a CPU, 512 MB of memory, a 1,024 MB volume, and that is attached to `network0`. Then allocate storage from `Pool-0001` for the volume. The volume is named `vm001_vol0`.

```
$ avcli vm-create --name vm001 --cpu 1 --memory 512 --
cdrom linux.iso --interfaces network0 \
--volumes 1024,Pool-0001,vm001_vol0
```

Create a VM named `vm001` with a CPU and 512 MB of memory, and that is attached to `network0` and `network1`. Create two volumes, where the first is 10 GB and the second is 50 GB. Allocate storage for these volumes from `Pool-0001` and `Pool-0002`, respectively.

```
$ avcli vm-create --name vm001 --cpu 1 --memory 512 --
cdrom linux.iso \
```

```
--interfaces network0 network1 \  
--volumes 10GB,Pool-0001 50GB,Pool-0002
```

Create a VM based on a kickstart template.

```
$ avcli vm-create --name vm001 --cpu 1 --memory 512 --kick-  
start template:o81 --interfaces network0 \  
--volumes 10GB
```

vm-delete

Usage

```
avcli vm-delete [--volumes volumes] [--wait] vm...
```

Description

The `vm-delete` command deletes the specified VMs, and optionally, the volumes attached to the VMs.

Options

<code>--volumes</code> <i>volumes</i>	Delete the volumes attached to the VM.
<code>--wait</code> <code>-w</code>	Wait for the command to complete.
<i>vm</i>	One or more VMs to be deleted.

Examples

```
avcli vm-delete vm1
```

```
avcli vm-delete --volumes vm1
```

```
avcli vm-delete --volumes vm1 vm2
```

vm-import**Usage**

```
avcli vm-import --archive filename.ova [--no-auto-start] [--cpu number] [--memory size] [--name vm-name] [--storage-groups groups] [--interfaces networks] [--volumes volumes] [--data] [--force] [--silent] [--dry-run] [--throttle] [--use-https]
```

Description

The `vm-import` command imports a VM from an OVA or OVF VM archive file.

Options

<code>--archive <i>filename.ova</i></code>	The OVA or OVF file archive to import.
<code>--no-auto-start</code>	Do not start the VM after the import has finished.
<code>--cpu <i>number</i></code>	The number of CPUs to assign to the VM. This defaults to the value in the archive.
<code>--memory <i>size</i></code>	The size of memory, in megabytes, to assign to the VM. This defaults to the value in the archive.
<code>--name <i>vm-name</i></code>	The name to assign to the VM. This defaults to the value in the archive.
<code>--storage-groups <i>groups</i></code>	The list of storage groups to use for allocating the VM's volumes. By default, all available storage groups are used. Allocation occurs in a round-robin fashion.
<code>--interfaces <i>networks</i></code>	The list of shared networks to assign to the VM's interfaces. By default, values in the archive or available shared networks are assigned.
<code>--volumes <i>volumes</i></code>	Import only these volumes. By default, all available volumes

	from the OVF are imported.
<code>--data</code>	Import data only for the specified volumes.
<code>--force</code>	When the OVF file is missing the <code>isBootable</code> flag (a known issue for Windows XP), assume that the VHD pointed to by OVF is the bootable one.
<code>--silent</code>	Suppress output.
<code>--dry-run</code>	Show the interface to the shared network and volume-to-storage-group assignments without actually importing or restoring a VM.
<code>--throttle</code>	Slow down the import/export operation. Valid values are: <ul style="list-style-type: none"> • <code>none</code>: Do not use throttling. This is the default value. • <code>low</code>: Slow down by about 25%. • <code>medium</code>: Slow down by about 50%. • <code>high</code>: Slow down by about 75%.
<code>--use-https</code>	Use secure HTTPS transport instead of the default streaming method (HTTP transport). Streaming over HTTPS provides slower performance than HTTP but is much more secure.

Examples

```
$ avcli vm-import --archive vm1.ova
$ avcli vm-import --archive vm1.ovf
$ avcli vm-import --name myVM --throttle low --archive vm1.ovf
$ avcli vm-import --cpu 2 --memory 1024 --archive vm1.ovf
$ avcli vm-import --interfaces network0 network1 --archive
vm1.ovf
```

```
$ avcli vm-import --storage-groups sm-0000 sm-0001 --archive  
vm1.ovf
```

```
$ avcli vm-import --volumes boot_vol vol3 --data vol3 --  
archive vm1.ovf
```

vm-info

Usage

```
avcli vm-info [vm...]
```

Description

The `vm-info` command displays information about all VMs, or optionally, about specific VMs.

Options

<i>vm</i>	One or more VMs for which to display information.
-----------	---------------------------------------------------

Examples

```
$ avcli vm-info
```

```
$ avcli vm-info vm1
```

```
$ avcli vm-info vm1 vm:o100
```

vm-migrate**Usage**

```
avcli vm-migrate [vm...] [--wait]
```

Description

The `vm-migrate` command migrates the specified VMs.

Options

<i>vm</i>	One or more VMs to migrate. Specify a VM either by name or ID.
<code>--wait</code> <code>-w</code>	Wait for the command to complete.

Examples

```
$ avcli vm-migrate vm1
```

```
$ avcli vm-migrate vm1 vm2
```

```
$ avcli vm-migrate vm1 vm:o100
```

vm-poweroff

Usage

```
avcli vm-poweroff [vm...] [--wait]
```

Description

The `vm-poweroff` command powers off the specified VMs.

Options

<code>vm</code>	One or more VMs to power off. Specify a VM either by name or ID.
<code>--wait</code> <code>-w</code>	Wait for the command to complete.

Examples

```
$ avcli vm-poweroff vm1  
$ avcli vm-poweroff vm1 vm2  
$ avcli vm-poweroff vm1 vm:o100
```

vm-poweron**Usage**

```
avcli vm-poweron [vm...] [--wait]
```

Description

The `vm-poweron` command powers on the specified VMs.

Options

<code>vm</code>	One or more VMs to power on. Specify a VM either by name or ID.
<code>--wait</code> <code>-w</code>	Wait for the command to complete.

Examples

```
$ avcli vm-poweron vm1  
$ avcli vm-poweron vm1 vm2  
$ avcli vm-poweron vm1 vm:o100
```

vm-reprovision

Usage

```
avcli vm-reprovision --name name [--cpu number] [--memory
size] [--addVolumes volumes] [--deleteVolumes volumes] [--keep-
Volumes volumes] [--interfaces networks]
```

Description

The `vm-reprovision` command reprovisions the specified VM.

Options

<code>--name <i>name</i></code>	Specify the VM to reprovision. Reprovision only one VM at a time. Specify a VM either by name or ID.
<code>--cpu <i>number</i></code>	The number of virtual CPUs. This defaults to the VM's current amount.
<code>--memory <i>size</i></code>	The size of memory, in megabytes. This defaults to the VM's current amount.
<code>--addVolumes <i>volumes</i></code>	<p>The list of volume(s) to create and attach to this VM. A <i>volume</i> is made up of three components, separated by commas:</p> <ul style="list-style-type: none"> • The size of the volume; required. • The storage-group name of the ID from which to carve storage • The volume name. <p>By default, the volume size is specified in megabytes, but you can use standard qualifiers such as KB, MB, GB, and TB.</p>
<code>--deleteVolumes <i>volumes</i></code>	The list of volume(s) that are currently attached to the specified VM to be deleted. Specify a volume either by name or ID.
<code>--keepVolumes <i>volumes</i></code>	The list of volume(s) that are currently attached to the specified

	VM to be kept attached to the VM. If you specify a volume that is currently attached but not specified in this list, the volume is detached (not destroyed) from the VM. Specify a volume either by name or ID.
<code>--interfaces <i>networks</i></code>	The list of networks to attach to the VM. Specify a network only once. The attached network must not be private.

Examples

```
$ avcli vm-reprovision --cpu 2 --name vm1
```

```
$ avcli vm-reprovision --cpu 2 --name vm:o100
```

```
$ avcli vm-reprovision --cpu 2 --memory 2048 --name vm:o100
```

Reprovision a VM named `vm001` that has a CPU, 512 MB of memory, a 1,024 MB volume and is attached to `network0`, and then allocate storage from `Pool-0001` for the volume. The volume is named `vm001_vol0`.

```
$ avcli vm-reprovision --cpu 1 --memory 512 --interfaces
network0 \
--addVolumes 1024,Pool-0001,vm001_vol0 --name vm1
```

Reprovision VM `vm1`, and then delete the volumes `volume:o411`, `data-vm1`, and `data-vm2` associated with it.

```
$ avcli vm-reprovision --deleteVolumes volume:o411 data-
vm1 data-vm2 --name vm1
```

Reprovision VM `vm1` with the new data volume `data-1-7`, delete volume `volume:o1043`, keep volumes `volume:o1`, `volume:o2`, `volume:o4`, and attach network interfaces `sharednetwork:o129` and `sharednetwork:o130`.

```
$ avcli vm-reprovision --cpu 3 --memory 3359 --addVolume
2500,storagegroup:o54,data-1-7 --deleteVolumes
volume:o1043 --keepVolumes volume:o1 volume:o2 volume:o4 -
--interfaces sharednetwork:o129 sharednetwork:o130
```

```
--name vm1
```

vm-shutdown

Usage

```
avcli vm-shutdown [vm...][--wait]
```

Description

The `vm-shutdown` command shuts down the specified VMs.

Options

<code>vm</code>	One or more VMs to shut down. Specify a VM either by name or ID.
<code>--wait</code> <code>-w</code>	Wait for the command to complete.

Examples

```
$ avcli vm-shutdown vm1
```

```
$ avcli vm-shutdown vm1 vm2
```

```
$ avcli vm-shutdown vm1 vm:o100
```

vm-unlock

Usage

```
avcli vm-unlock [vm...]
```

Description

The `vm-unlock` command unlocks the specified VMs. For example, VM import operations set the lock to prevent a VM from being started or modified while the operation is occurring. If an operation fails unexpectedly, leaving the VM locked, use this command to unlock that VM.

Options

<i>vm</i>	One or more VMs to unlock. Specify a VM either by name or ID.
-----------	---------------------------------------------------------------

Examples

```
$ avcli vm-unlock vm1
```

```
$ avcli vm-unlock vm:o100
```

volume-info**Usage**

```
avcli volume-info [volume...]
```

Description

The `volume-info` command displays information about all volumes, or optionally, only about those specified.

Options

<i>volume</i>	A volume about which to display information.
---------------	----------------------------------------------

12

Chapter 12: System Reference Information

See the following topics for reference information

- ["Compatible Guest Operating Systems" on page 316](#)
- ["Physical Machine System Requirements" on page 317](#)
- ["Important Physical Machine and Virtual Machine Considerations" on page 319](#)

Compatible Guest Operating Systems

The following are compatible as guest operating systems for virtual machines (VMs) running on everRun Enterprise systems.

Vendor	Versions
Microsoft Windows Desktop	Windows 7, 32-bit, 64-bit Windows 8, 64-bit Enterprise Windows 8.1, 64-bit Enterprise
Windows Server 2008	32-bit, SP2 64-bit, SP2, R2 SP1 Web, Small Business, Standard, Enterprise, Datacenter
Microsoft Windows Small Business Server 2011	64-bit (only OS) Standard, Essential, Premium Add-On

Vendor	Versions
Microsoft Windows Server 2012	64-bit (only OS) Foundation, Essentials, Standard, Datacenter
Microsoft Windows Server 2012 R2	64-bit (only OS) Foundation, Essentials, Standard, Datacenter
Red Hat Enterprise Linux 6	Red Hat 6.4, 6.5, 64-bit Workstation, Server
CentOS 6	CentOS 6.4, 6.5, 64-bit
SUSE Linux Enterprise Server	SLES 11, SP3 64-bit
Ubuntu	12.04 13.1, 64-bit

Physical Machine System Requirements

The following table presents the minimum and maximum capacities of the listed devices for physical machines running on everRun Enterprise systems.

Physical Device	Min-imum	Max-imum Tested	Architected	Notes
CPUs: Intel® Xeon® E3-XXXX Processor Intel Xeon E3-XXXX v2 Processor Intel Xeon E5-XXXX Processor Intel Xeon E5-XXXX v2 Processor	1	2	No Practical Limit	

Physical Device	Min-imum	Max-imum Tested	Architected	Notes
Number CPU Sockets per Physical Machine	1	2	No Practical Limit	
Physical Memory	8 GB	384 GB	No Practical Limit	
Internal Disk Count per Physical Machine	2	24	No Practical Limit	Minimum 2 drives per PM for FT. VM disks/volumes are replicated on both PMs.
Total Disk Capacity	36 GB	9.4 TB	No Limit	
Logical Volume Size (everRun Boot Volumes)		5 TB	No Limit	
Logical Data Volume Size (Non-boot Volumes)		5 TB	No Limit	
Management ENET Ports	1	1	1	1 per system required.
A-Link ENET Ports	1 on each PM	8 on each PM		2 recommended. No VM can have more than 2. Maximum of 8 (for 4 or more guests)
Business ENET Ports	1	20		Can be shared with Management

Physical Device	Min-imum	Max-imum Tested	Architected	Notes
				link.
Quorum Servers	0	2		

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:

- ["Physical Machine System Requirements" on page 317](#)
- ["Virtual Machine Recommendations and Limits" on page 319](#)
- ["Combined Virtual Machine Maximums" on page 321](#)
- ["Important Considerations" on page 321](#)

Virtual Machine Recommendations and Limits

Virtual machines (VMs) require certain [CPU core resources](#) and have [other limits](#) for memory, networks, and storage.

Recommended Number of CPU Cores

The number of cores recommended for everRun Enterprise workloads depends upon the number of vCPUs in each VM and the types of the VMs as described below:

Item	Number of Physical Threads
Fixed system overhead (host and system management)	2
Each FT guest with n vCPUs	$n + 2$ (typical)
Each HA guest with n vCPUs	$n + 1$ (typical)



Note: A non-hyperthreaded physical CPU core can handle 1 physical thread. A hyper-threaded physical CPU core can handle 2 physical threads.

The actual number of required threads depends upon the workload. The guidelines above should cover most workloads. Since any given workload may require fewer or more threads, it's good practice to test and characterize your specific workload.

Examples

A single 4-vCPU FT guest typically requires:

- 2 threads for host/system management
- 6 threads for guest
 - **8 threads total** (a single socket 4-core hyper-threaded system)

Four 5-vCPU FT guests typically require:

- 2 threads for host/system management
- 7 threads for first guest
- 7 threads for second guest
- 7 threads for third guest
- 7 threads for fourth guest
 - **30 threads total** (a dual socket 8-core hyper-threaded system)

Virtual Machine Limits

For systems with many or large virtual machines (VMs), configure everRun with 10 Gb sync links, and for the everRun software itself, 4 vCPUs and 4096 MB. Refer to the **Preferences -> Systems Resources** page in the everRun Availability Console for instructions on setting the everRun Enterprise system resources to the maximum.

The following table lists an everRun Enterprise system's maximum VM capacities.

Item	Maximum
Max. VCPUs per FT VM	8

Item	Maximum
Max. VCPUs per HA VM	16
Max. Memory per FT VM	256 GB
Max. Memory per HA VM	256 GB
Max. Availability Links per VM	2
Max. Virtual Networks per VM	20
Max. Storage Volumes per VM	12
Max. Volume Size	Less than 2 TB

Combined Virtual Machine Maximums

The following table presents the combined maximums of virtual machines (VMs) and virtual NICs that can run on everRun Enterprise systems.

Virtual Device	Maximum
Total FT VMs	4
Total VMs (FT and HA combined)	12
Total Virtual Network Interface cards (NICs)	20

Important Considerations

Note the following important considerations.

Feature	Comment
everRunSystem Disk	Recommended minimum configuration for physical machines: <ul style="list-style-type: none"> One logical volume protected by RAID1, RAID 5, RAID 6, or RAID 10

Feature	Comment
	<p>or</p> <ul style="list-style-type: none"> Two non-RAID or RAID 0 volumes. <p>When using multiple volumes per RAID set, the RAID set should set a type that provides redundancy, such as RAID1, RAID5, or RAID10.</p>
USB CD/DVD Drive	USB CD/DVD drives are supported on all platforms for everRun installation.
Direct-Attach Tape Drives	Access to direct-attach tape drives by the guests is not supported. Stratus recommends using network-attach tape drives.
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.
SSD Support	everRun supports solid state drives according to the storage controller vendor specifications.
System Management	everRun Enterprise system management does not work at run-level 5.

13

Chapter 13: 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 everRun Enterprise system to use SNMP, see ["Configuring SNMP Settings" on page 75](#).

To view the contents of the MIB file, see ["MIB File Contents" on page 324](#)

MIB File Contents

The management information base (MIB) is a file that describes the set of network objects that Simple Network Management Protocol (SNMP) can manage on an everRun system.

The format of the MIB is defined as part of the SNMP.

The MIB file is shown in its entirety below.

```
-- =====  
--  
--  COPYRIGHT (c) 2001 - 2014  Stratus Technologies Bermuda Ltd.  
--  All Rights Reserved.  
--  
-- =====  
  
-- =====  
--  
--  @File:  
--    STRATUS-EVERRUN-MIB.txt  
--
```

everRun Enterprise User's Guide

```
-- @Revision:
--   2.0
--
-- @Description:
--   This file defines the Stratus everRun SNMP MIB.
--   Definitions for everRun agents appear here.
--   Stratus MIB definitions for other agents are not in this file.
--
-- =====

STRATUS-MIB DEFINITIONS ::= BEGIN

IMPORTS
    enterprises          FROM RFC1155-SMI
    OBJECT-TYPE          FROM RFC-1212
    DisplayString        FROM RFC1213-MIB
    TRAP-TYPE            FROM RFC-1215;

Boolean ::= INTEGER {
    unknown(1),
    false(2),
    true(3)
}

ToggleState ::= INTEGER {
    enabled(1),
    disabled(2)
}

-- This data type is to indicate true or false.

-- =====
-- Stratus Enterprise tree structure
-- =====
-- stratus enterprise : 1.3.6.1.4.1.458
--
stratus          OBJECT IDENTIFIER ::= { enterprises 458 }

-- =====
```

```
-- Major categories under the Stratus namespace.
-- Note: Values less than 101 are not used to prevent collision with
-- old products.
-- =====
experimental    OBJECT IDENTIFIER ::= { stratus 101 }
agentInfo       OBJECT IDENTIFIER ::= { stratus 102 }
systemInfo      OBJECT IDENTIFIER ::= { stratus 103 }
productIdent    OBJECT IDENTIFIER ::= { stratus 104 }
ftServerOid    OBJECT IDENTIFIER ::= { stratus 105 }
stcpOid         OBJECT IDENTIFIER ::= { stratus 106 }
ftLinuxOid     OBJECT IDENTIFIER ::= { stratus 107 }
avanceOid       OBJECT IDENTIFIER ::= { stratus 110 }
everRunOid     OBJECT IDENTIFIER ::= { stratus 115 }

-- =====
-- The Agent Information table is used to provide information about
-- the capabilities of the SNMP agent.
-- =====
sraAgentMibFamily OBJECT-TYPE
    SYNTAX      INTEGER {
                    stcp(1),
                    ftServer(2),
                    ftlinux(3),
                    avance(4),
                    everRun(5)
                }
    ACCESS      read-only
    STATUS      mandatory
    DESCRIPTION
        "This variable indicates which OIDs are supported by the agent.
        When support for variables and/or traps are removed from an
        agent, a new family must be created."
    ::= { agentInfo 1 }

sraAgentMibRevision OBJECT-TYPE
    SYNTAX      INTEGER {
                    rev01(1)
                }
    ACCESS      read-only
    STATUS      mandatory
```

DESCRIPTION

"This variable indicates whether variables and/or traps have been added to the MIB. When a MIB family is created this is initially one. When OIDs are added to those an agent supports, this integer is incremented. Each time a MIB is published, the corresponding Revision will be defined in the MIB."

::= { agentInfo 2 }

-- =====
-- The System Information table provides information about system as a
-- whole. These variables are platform independent.
-- =====

sraSiSystemType OBJECT-TYPE

SYNTAX OBJECT IDENTIFIER
ACCESS read-only
STATUS mandatory

DESCRIPTION

"The authoritative identification of the hardware and software in the entity. This value provides an easy and unambiguous means for determining 'what kind of box' is being managed. This value is an OID that indicates the product family, operating system and CPU architecture. Values are enumerated in the Product Identification (OID 104) table."

::= { systemInfo 1 }

sraSiManufacturer OBJECT-TYPE

SYNTAX DisplayString
ACCESS read-only
STATUS mandatory

DESCRIPTION

"This value is a string to indicate the manufacturer of the system. If unknown, the agent may return a null string."

::= { systemInfo 2 }

sraSiModel OBJECT-TYPE

SYNTAX DisplayString
ACCESS read-only
STATUS mandatory

DESCRIPTION

"This value is a string to indicate the model of the system."

```
        If unsupported the agent may return a null string."
 ::= { systemInfo 3 }

sraSiOverallSystemStatus OBJECT-TYPE
    SYNTAX  INTEGER  {
        unsupported(1),
        noFaults(2),
        systemFault(3),
        systemDown(4)
    }
    ACCESS   read-only
    STATUS   mandatory
    DESCRIPTION
        "This integer indicates the overall status of the system."
 ::= { systemInfo 4 }

sraSiSystemName OBJECT-TYPE
    SYNTAX   DisplayString
    ACCESS   read-only
    STATUS   mandatory
    DESCRIPTION
        "This value is a string representing the network name of the
        system.  This is expected to be unique on a LAN but possibly
        not globally unique.  If unsupported by the agent, a null
        string may be returned.  When the OS is Windows, this is the
        *computer name* portion of the network id, or the Lan Manager
        name of the computer (e.g. PCAT).  In contrast, the MIB-II
        sysName is typically the fully-qualified domain name
        (e.g. pcat.mno.stratus.com).  On VOS, this is the system and
        module name (e.g. %sys#m1).  On UNIX and Linux this is the
        hostname."
 ::= { systemInfo 5 }

sraSiSystemSerialNumber OBJECT-TYPE
    SYNTAX   DisplayString
    ACCESS   read-only
    STATUS   mandatory
    DESCRIPTION
        "This value is a string containing the serial number of the
        system.  If unsupported by the agent, a null string may be
        returned."
```

everRun Enterprise User's Guide

```
 ::= { systemInfo 6 }

sraSiSiteID OBJECT-TYPE
    SYNTAX      DisplayString
    ACCESS      read-only
    STATUS      mandatory
    DESCRIPTION
        "This string value contains the SiteID. SiteID is part of
        the RSN/ASN service model."
 ::= { systemInfo 7 }

sraSiCpuFamily OBJECT-TYPE
    SYNTAX      INTEGER {
                unsupported(1),
                m68k(2),
                i860(3),
                hppa(4),
                ia32(5),
                ia64(6)
                }
    ACCESS      read-only
    STATUS      mandatory
    DESCRIPTION
        "This value is an integer that indicates the CPU architecture."
 ::= { systemInfo 8 }

sraSiOsType OBJECT-TYPE
    SYNTAX      INTEGER {
                unsupported(1),
                ftx(2),
                hpux(3),
                ftlinux(4),
                vos(5),
                windows(6),
                avance(7),
                everRun(8)
                }
    ACCESS      read-only
    STATUS      mandatory
    DESCRIPTION
        "This value is an integer that indicates Operating System type."
```

```

 ::= { systemInfo 9 }

-- =====
-- The Product Identification table is used to identify specific Stratus
-- products. This table defines OIDs but there are no variables. Where
-- possible these will be used as the value of the RFC-1213 MIB-II
-- system.sysObjectID variable. However, with a non-Stratus OS, like
-- ftLinux and Windows, MIB-II system.sysObjectID is not under our control.
-- Consequently these same values are reported in the Stratus variable
-- sraSiSystemType.
-- =====

osFTX                OBJECT IDENTIFIER ::= { productIdent 1 }
sraProductIdFtxJetta OBJECT IDENTIFIER ::= { osFTX 1 }
sraProductIdFtxPolo  OBJECT IDENTIFIER ::= { osFTX 2 }

osHPUX               OBJECT IDENTIFIER ::= { productIdent 2 }
sraProductIdHpuxPolo OBJECT IDENTIFIER ::= { osHPUX 1 }

osftLinux            OBJECT IDENTIFIER ::= { productIdent 3 }
sraProductIdLnxFtsIa32 OBJECT IDENTIFIER ::= { osftLinux 1 }

osVOS                OBJECT IDENTIFIER ::= { productIdent 4 }
sraProductIdVos68k   OBJECT IDENTIFIER ::= { osVOS 1 }
sraProductIdVosI860  OBJECT IDENTIFIER ::= { osVOS 2 }
sraProductIdVosJetta OBJECT IDENTIFIER ::= { osVOS 3 }
sraProductIdVosIa32  OBJECT IDENTIFIER ::= { osVOS 4 }

osWindowsFt         OBJECT IDENTIFIER ::= { productIdent 5 }
sraProductIdWinFtsIa32 OBJECT IDENTIFIER ::= { osWindowsFt 1 }
sraProductIdWinFtsIa64 OBJECT IDENTIFIER ::= { osWindowsFt 2 }

osRadio              OBJECT IDENTIFIER ::= { productIdent 6 }
sraProductIdWinRadIa32 OBJECT IDENTIFIER ::= { osRadio 1 }

osAvance             OBJECT IDENTIFIER ::= { productIdent 10 }
sraProductIdAvance   OBJECT IDENTIFIER ::= { osAvance 1 }

osEverRun            OBJECT IDENTIFIER ::= { productIdent 15 }
sraProductIdEverRun  OBJECT IDENTIFIER ::= { osEverRun 1 }

```

everRun Enterprise User's Guide

```
-- =====
-- The following table contains OIDs unique to the everRun MIB.
-- There are three groups of OIDs within this table:
-- OIDs that identify GET/SET variables,
-- OIDs that identify everRun TRAPs, and
-- OIDs used to identify variable fields in TRAP PDUs.
-- =====

everRunVar          OBJECT IDENTIFIER ::= { everRunOid 1 }
everRunTrapId      OBJECT IDENTIFIER ::= { everRunOid 2 }
everRunTrapData    OBJECT IDENTIFIER ::= { everRunOid 3 }

-- everRun GET/SET variables

everRunAvailableVirtualMemory OBJECT-TYPE
    SYNTAX      INTEGER
    ACCESS      read-only
    STATUS      mandatory
    DESCRIPTION
        "This integer contains the available virtual memory of the system
        in gigabytes."
    ::= { everRunVar 1 }

everRunVirtualCPUsTotal OBJECT-TYPE
    SYNTAX      INTEGER
    ACCESS      read-only
    STATUS      mandatory
    DESCRIPTION
        "This integer contains the total number of virtual CPUs on the system."
    ::= { everRunVar 2 }

everRunVirtualCPUsInUse OBJECT-TYPE
    SYNTAX      INTEGER
    ACCESS      read-only
    STATUS      mandatory
    DESCRIPTION
        "This integer contains the number of virtual CPUs currently in use on
        the system."
```

```
 ::= { everRunVar 3 }

everRunVirtualCPUsMaxPerVM OBJECT-TYPE
    SYNTAX      INTEGER
    ACCESS      read-only
    STATUS      mandatory
    DESCRIPTION
        "This integer contains the maximum number of virtual CPUs that can be
        assigned to a virtual machine."
 ::= { everRunVar 4 }

-- everRunVirtualCPUsPercentageUsed OBJECT-TYPE
--     SYNTAX      INTEGER
--     ACCESS      read-only
--     STATUS      mandatory
--     DESCRIPTION
--         "This integer contains the percentage of available virtual CPU capacity
--         that is in use on the system."
--     ::= { everRunVar 5 }

everRunStorageTotal OBJECT-TYPE
    SYNTAX      INTEGER
    ACCESS      read-only
    STATUS      mandatory
    DESCRIPTION
        "This integer contains the total amount of storage on the system in gigabytes."
 ::= { everRunVar 5 }

everRunStorageUsed OBJECT-TYPE
    SYNTAX      INTEGER
    ACCESS      read-only
    STATUS      mandatory
    DESCRIPTION
        "This integer contains the amount of storage in use on the system in gigabytes."
 ::= { everRunVar 6 }

-- everRunStorageUsedByManagement OBJECT-TYPE
--     SYNTAX      INTEGER
--     ACCESS      read-only
--     STATUS      mandatory
--     DESCRIPTION
```

everRun Enterprise User's Guide

```
--          "This integer contains the amount of storage in use by management on the system
--          in gigabytes."
--      ::= { everRunVar 7 }

everRunStorageFree OBJECT-TYPE
    SYNTAX      INTEGER
    ACCESS      read-only
    STATUS      mandatory
    DESCRIPTION
        "This integer contains the amount of unused storage on the system in gigabytes."
    ::= { everRunVar 7 }

-- everRunDiskReadBytes OBJECT-TYPE
--     SYNTAX      INTEGER
--     ACCESS      read-only
--     STATUS      mandatory
--     DESCRIPTION
--         "This integer contains the percentage of available disk throughput on the system that
--         is being consumed by disk reads."
--     ::= { everRunVar 10 }

-- everRunDiskWriteBytes OBJECT-TYPE
--     SYNTAX      INTEGER
--     ACCESS      read-only
--     STATUS      mandatory
--     DESCRIPTION
--         "This integer contains the percentage of available disk throughput on the system that
--         is being consumed by disk writes."
--     ::= { everRunVar 11 }

everRunIPAddress OBJECT-TYPE
    SYNTAX      IpAddress
    ACCESS      read-only
    STATUS      mandatory
    DESCRIPTION
        "This IP address is the IP address of the system.  It corresponds to the
        fully qualified domain name of the system."
    ::= { everRunVar 8 }

-- everRunNetworkReadBytes OBJECT-TYPE
--     SYNTAX      INTEGER
```

```
-- ACCESS      read-only
-- STATUS      mandatory
-- DESCRIPTION
--             "This integer contains the percentage of available network bandwidth on the system
that
--             is being consumed by network reads."
-- ::= { everRunVar 13 }

-- everRunNetworkWriteBytes OBJECT-TYPE
-- SYNTAX      INTEGER
-- ACCESS      read-only
-- STATUS      mandatory
-- DESCRIPTION
--             "This integer contains the percentage of available network bandwidth on the system
that
--             is being consumed by network writes."
-- ::= { everRunVar 14 }

everRunAlertNumber OBJECT-TYPE
    SYNTAX      INTEGER
    ACCESS      read-only
    STATUS      mandatory
    DESCRIPTION
        "This integer contains the number of entries in the everRunAlertTable table."
    ::= { everRunVar 9 }

everRunAlertTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF everRunAlertEntry
    ACCESS      read-only
    STATUS      optional
    DESCRIPTION
        "This table contains an entry for each alert log that has been generated on this sys-
tem."
    ::= { everRunVar 10 }

everRunAlertEntry OBJECT-TYPE
    SYNTAX      everRunAlertEntry
    ACCESS      read-only
    STATUS      optional
    DESCRIPTION
        "This entry represents one alert in the everRunAlertTable table."
```

everRun Enterprise User's Guide

```
INDEX      { everRunAlertIndex }
 ::= { everRunAlertTable 1 }

everRunAlertEntry ::= SEQUENCE {
    everRunAlertIndex      INTEGER,
    everRunAlertSeverity   INTEGER,
    everRunAlertType       INTEGER,
    everRunAlertSource     DisplayString,
    everRunAlertDateTime   DisplayString,
    everRunAlertCallHomeSent Boolean,
    everRunAlertEAlertSent Boolean,
    everRunAlertSNMPTrapSent Boolean,
    everRunAlertInformation DisplayString,
    everRunAlertSNMPTrapOID OBJECT IDENTIFIER }

everRunAlertIndex OBJECT-TYPE
    SYNTAX      INTEGER(0..65535)
    ACCESS      read-only
    STATUS      optional
    DESCRIPTION
        "This index value uniquely identifies the alert represented by this entry."
    ::= { everRunAlertEntry 1 }

everRunAlertSeverity OBJECT-TYPE
    SYNTAX      INTEGER {
        clear(0),
        informational(1),
        minor(2),
        major(3),
        serious(4),
        critical(5)
    }
    ACCESS      read-only
    STATUS      optional
    DESCRIPTION
        "This value represents the severity of the alert."
    ::= { everRunAlertEntry 2 }

everRunAlertType OBJECT-TYPE
    SYNTAX      DisplayString
    ACCESS      read-only
```

```

STATUS    optional
DESCRIPTION
    "This value represents the type of the alert."
 ::= { everRunAlertEntry 3 }

everRunAlertSource  OBJECT-TYPE
    SYNTAX      DisplayString
    ACCESS      read-only
    STATUS      optional
    DESCRIPTION
        "This string contains the source of the alert.  This could be a device or a node."
 ::= { everRunAlertEntry 4 }

everRunAlertDateTime  OBJECT-TYPE
    SYNTAX      DisplayString
    ACCESS      read-only
    STATUS      optional
    DESCRIPTION
        "This string contains the date and time that the alert was generated."
 ::= { everRunAlertEntry 5 }

everRunAlertCallHomeSent  OBJECT-TYPE
    SYNTAX      Boolean
    ACCESS      read-only
    STATUS      optional
    DESCRIPTION
        "This boolean value indicates whether or not a CallHome message was sent for this
alert."
 ::= { everRunAlertEntry 6 }

everRunAlertEAlertSent  OBJECT-TYPE
    SYNTAX      Boolean
    ACCESS      read-only
    STATUS      optional
    DESCRIPTION
        "This boolean value indicates whether or not an eAlert was sent for this alert."
 ::= { everRunAlertEntry 7 }

everRunAlertSNMPTrapSent  OBJECT-TYPE
    SYNTAX      Boolean
    ACCESS      read-only

```

everRun Enterprise User's Guide

STATUS optional

DESCRIPTION

"This boolean value indicates whether or not an SNMP trap was sent for this alert."

::= { everRunAlertEntry 8 }

everRunAlertInformation OBJECT-TYPE

SYNTAX DisplayString

ACCESS read-only

STATUS optional

DESCRIPTION

"This string contains explanatory text regarding the alert. This can include more details regarding

the cause of the alert and the device/node that caused the alert to be generated."

::= { everRunAlertEntry 9 }

everRunAlertSNMPTrapOID OBJECT-TYPE

SYNTAX OBJECT IDENTIFIER

ACCESS read-only

STATUS optional

DESCRIPTION

"This string contains the OID of the trap associated with this alert. Even if the trap is not sent, this field will contain the OID of the trap that would have been sent."

::= { everRunAlertEntry 10 }

everRunAuditNumber OBJECT-TYPE

SYNTAX INTEGER

ACCESS read-only

STATUS mandatory

DESCRIPTION

"This integer contains the number of entries in the everRunAuditTable table."

::= { everRunVar 11 }

everRunAuditTable OBJECT-TYPE

SYNTAX SEQUENCE OF everRunAuditEntry

ACCESS read-only

STATUS optional

DESCRIPTION

"This table contains an entry for each audit that has been generated on this system."

::= { everRunVar 12 }

everRunAuditEntry OBJECT-TYPE

```

SYNTAX      everRunAuditEntry
ACCESS      read-only
STATUS      optional
DESCRIPTION
            "This entry represents one audit in the everRunAuditTable table."
INDEX       { everRunAuditIndex }
 ::= { everRunAuditTable 1 }

everRunAuditEntry ::= SEQUENCE {
    everRunAuditIndex      INTEGER,
    everRunAuditDateTime   DisplayString,
    everRunAuditUsername   DisplayString,
    everRunAuditOriginatingHost IpAddress,
    everRunAuditAction     DisplayString
}

everRunAuditIndex  OBJECT-TYPE
    SYNTAX      INTEGER(0..65535)
    ACCESS      read-only
    STATUS      optional
    DESCRIPTION
        "This index value uniquely identifies the audit represented by this entry."
 ::= { everRunAuditEntry 1 }

everRunAuditDateTime  OBJECT-TYPE
    SYNTAX      DisplayString
    ACCESS      read-only
    STATUS      optional
    DESCRIPTION
        "This string contains the date and time that the audit was generated."
 ::= { everRunAuditEntry 2 }

everRunAuditUsername  OBJECT-TYPE
    SYNTAX      DisplayString
    ACCESS      read-only
    STATUS      optional
    DESCRIPTION
        "This string contains the username of the user that caused the audit to be generated."
 ::= { everRunAuditEntry 3 }

everRunAuditOriginatingHost  OBJECT-TYPE

```

everRun Enterprise User's Guide

```
SYNTAX      IPAddress
ACCESS      read-only
STATUS      optional
DESCRIPTION
            "This is the address of the host that originated the audit."
::= { everRunAuditEntry 4 }

everRunAuditAction  OBJECT-TYPE
SYNTAX      DisplayString
ACCESS      read-only
STATUS      optional
DESCRIPTION
            "This string contains a description of the action being audited."
::= { everRunAuditEntry 5 }

-- everRun TRAP PDU Data Fields
-- This table contains variables that may be included in trap PDUs.

everRunTrapDescription  OBJECT-TYPE
SYNTAX      DisplayString
ACCESS      read-only
STATUS      mandatory
DESCRIPTION
            "This string contains descriptive data -- suitable for display -- about the trap."
::= { everRunTrapData 1 }

--everRunTrapObject  OBJECT-TYPE
-- SYNTAX      OBJECT IDENTIFIER
-- ACCESS      read-only
-- STATUS      mandatory
-- DESCRIPTION
--            "This OID represents the object for which the trap is concerned."

-- everRun Traps
--
-- All everRun traps use *everRunTrapId* as the enterprise OID.
-- The traps are distinguished by a unique enterprise-specific TrapId.
-- The TrapId is the last token, following ::= in the TRAP-TYPE macro
-- invocation.
```

```
--  
  
everRunGenericTrap TRAP-TYPE  
    ENTERPRISE everRunTrapId  
    VARIABLES {  
        everRunTrapDescription  
    }  
    DESCRIPTION  
    "Generic Trap."  
    ::= 1  
  
everRunGuestCrashedTrap TRAP-TYPE  
    ENTERPRISE everRunTrapId  
    VARIABLES {  
        everRunTrapDescription  
    }  
    DESCRIPTION  
    "Guest Crashed Trap."  
    ::= 2  
  
everRunNodeUnreachableTrap TRAP-TYPE  
    ENTERPRISE everRunTrapId  
    VARIABLES {  
        everRunTrapDescription  
    }  
    DESCRIPTION  
    "Node Unreachable Trap."  
    ::= 3  
  
everRunNodeMaintenanceTrap TRAP-TYPE  
    ENTERPRISE everRunTrapId  
    VARIABLES {  
        everRunTrapDescription  
    }  
    DESCRIPTION  
    "Node Maintenance Trap."  
    ::= 4  
  
everRunDoubleFaultPredictionTrap TRAP-TYPE  
    ENTERPRISE everRunTrapId  
    VARIABLES {
```

everRun Enterprise User's Guide

```
        everRunTrapDescription
    }
DESCRIPTION
"Double Fault Prediction Trap."
 ::= 5

everRunPredictFaultOnSingleSystemNodeTrap TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"Predict Fault On Single System Node Trap."
 ::= 6

everRunDiskProblemTrap TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"Disk Problem Trap."
 ::= 7

everRunDetectionOfBadNetworkTrap TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"Detection Of Bad Network Trap."
 ::= 8

everRunDetectionOfBadSensorOnChassisTrap TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"Detection Of Bad Sensor On Chassis Trap."
 ::= 9
```

```
everRunNodeRebootedUnexpectedlyTrap TRAP-TYPE
    ENTERPRISE  everRunTrapId
    VARIABLES   {
                everRunTrapDescription
                }
    DESCRIPTION
    "Node Rebooted Unexpectedly Trap."
    ::= 10
```

```
everRunNodeBlacklistTrap TRAP-TYPE
    ENTERPRISE  everRunTrapId
    VARIABLES   {
                everRunTrapDescription
                }
    DESCRIPTION
    "Node Blacklist Trap."
    ::= 11
```

```
everRunVMBlacklistedTrap TRAP-TYPE
    ENTERPRISE  everRunTrapId
    VARIABLES   {
                everRunTrapDescription
                }
    DESCRIPTION
    "VM Blacklisted Trap."
    ::= 12
```

```
everRunVMBootFailedTrap TRAP-TYPE
    ENTERPRISE  everRunTrapId
    VARIABLES   {
                everRunTrapDescription
                }
    DESCRIPTION
    "VM Boot Failed Trap."
    ::= 13
```

```
unitPredictFaultOnSingleNodeUnit TRAP-TYPE
    ENTERPRISE  everRunTrapId
    VARIABLES   {
```

everRun Enterprise User's Guide

```
        everRunTrapDescription
    }
DESCRIPTION
    "Predict Fault On Single System Node."
 ::= 20

unitNoQuorum TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
DESCRIPTION
    "Cannot establish quorum."
 ::= 21

unitCallHomeNotEnabled TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
DESCRIPTION
    "Call Home Not Enabled."
 ::= 22

unitDialInNotEnabled TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
DESCRIPTION
    "Dial-In Not Enabled."
 ::= 23

unitEAlertNotEnabled TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
DESCRIPTION
    "E-Alert Notification Not Enabled."
 ::= 24
```

```
unitSnmpTrapNotEnabled TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "SNMP Trap Notification Not Enabled."
    ::= 25
```

```
unitNtpNotEnabled TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "NTP Time Synchronization Not Enabled."
    ::= 26
```

```
vmBlacklist TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "VMBlacklist."
    ::= 27
```

```
vmCrashed TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "Guest Crashed."
    ::= 28
```

```
vmBootFailed TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
```

everRun Enterprise User's Guide

```
        }
    DESCRIPTION
    "VMBootFailed."
    ::= 29

nodeUnreachable TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "Node Unreachable."
    ::= 32

nodeUnexpectedlyOff TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "Node Unreachable."
    ::= 33

nodeFailed TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "Node Unreachable."
    ::= 34

nodeBlacklist TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "NodeBlacklist."
    ::= 35
```

```
nodeMaintenance TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "Node Maintenance."
    ::= 36
```

```
nodeUnexpectedRebooted TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "Node rebooted unexpectedly."
    ::= 37
```

```
nodeVmxNotEnabled TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "PM Does Not Have VMX Enabled."
    ::= 38
```

```
nodeNxMismatch TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "PM Setting For NX Mismatch."
    ::= 39
```

```
nodeBootOrderIsIncorrect TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
```

everRun Enterprise User's Guide

```
DESCRIPTION
"PM Boot Order Is Incorrect."
::= 40

nodeOldSoftwareVersionFault TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"Node requires upgrade."
::= 41

nodeRunningOnBattery TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"On Battery."
::= 44

nodeRunningOnLowBattery TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"Low Battery - PM Shutdown."
::= 45

nodeLastNodeRunningOnLowBattery TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"Low Battery - Unit Shutdown."
::= 46

nodeExiled TRAP-TYPE
```

```
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"Node Exiled."
 ::= 47

diskFailed TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"Disk problem."
 ::= 48

diskNotPresent TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"Disk problem."
 ::= 49

diskIsMissing TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"PM Missing a Required Disk."
 ::= 50

diskIsTooSmall TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
```

everRun Enterprise User's Guide

```
"PM Disk is Too Small."
::= 51

nodeSingleDiskNotRedundant TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "PM System Disk is Not Redundant."
    ::= 52

networkNoLink TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "Detection of Bad Network."
    ::= 53

networkFailedPort TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "Detection of Bad Network."
    ::= 54

networkBadConnectivity TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "Detection of Bad Network."
    ::= 55

networkSlowBusiness TRAP-TYPE
    ENTERPRISE everRunTrapId
```

```
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"Detection of Slow Business Network."
 ::= 56

networkSlowPrivate TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"Detection of Slow Private Network."
 ::= 57

networkIsMissing TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"PM Does Not Have a Required Local Network."
 ::= 58

pdiskBroken TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"Physical Disk Problem."
 ::= 59

pdiskNotPresent TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"Physical Disk Problem."
```

everRun Enterprise User's Guide

```
 ::= 60

pdiskForeign TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "Physical Disk Problem."
 ::= 61

pdiskPredictFault TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "Physical Disk Problem."
 ::= 62

sensorMinor TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "Detection of Bad Sensor on chassis."
 ::= 63

sensorModerate TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "Detection of Bad Sensor on chassis."
 ::= 64

controllerBasicSupport TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
```

```
        everRunTrapDescription
    }
DESCRIPTION
    "Disk controller is not fully supported."
 ::= 67

nodePmModelNotSupported TRAP-TYPE
    ENTERPRISE  everRunTrapId
    VARIABLES   {
        everRunTrapDescription
    }
DESCRIPTION
    "PM is not a supported model."
 ::= 68

nodeSystemStorageNotRedundant TRAP-TYPE
    ENTERPRISE  everRunTrapId
    VARIABLES   {
        everRunTrapDescription
    }
DESCRIPTION
    "System Storage Not Redundant."
 ::= 69

unitProcIncompatVAPICSecondaryExec TRAP-TYPE
    ENTERPRISE  everRunTrapId
    VARIABLES   {
        everRunTrapDescription
    }
DESCRIPTION
    "Processor Incompatibility - Secondary Exec Virtual APIC Access."
 ::= 70

unitWarningSwap TRAP-TYPE
    ENTERPRISE  everRunTrapId
    VARIABLES   {
        everRunTrapDescription
    }
DESCRIPTION
    "Warning Swap."
 ::= 74
```

everRun Enterprise User's Guide

```
unitFatalSwap TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "Fatal Swap."
    ::= 75
```

```
unitSinglePM TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "Single PM Detected."
    ::= 77
```

```
unitEalertFailure TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "E-Alert Failure Detected."
    ::= 78
```

```
unitLicenseExpired TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "License Expired."
    ::= 79
```

```
unitLicenseAboutToExpire TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
```

```
        }  
DESCRIPTION  
    "License About to Expire."  
 ::= 80  
  
unitSnmpTrapFailure TRAP-TYPE  
    ENTERPRISE  everRunTrapId  
    VARIABLES   {  
                everRunTrapDescription  
    }  
DESCRIPTION  
    "SNMP Trap Failure Detected."  
 ::= 81  
  
unitCallHomeFailure TRAP-TYPE  
    ENTERPRISE  everRunTrapId  
    VARIABLES   {  
                everRunTrapDescription  
    }  
DESCRIPTION  
    "Call-Home Failure Detected."  
 ::= 82  
  
controllerRAIDBatteryFailed TRAP-TYPE  
    ENTERPRISE  everRunTrapId  
    VARIABLES   {  
                everRunTrapDescription  
    }  
DESCRIPTION  
    "RAID Battery Failed."  
 ::= 83  
  
controllerRAIDBatteryMissing TRAP-TYPE  
    ENTERPRISE  everRunTrapId  
    VARIABLES   {  
                everRunTrapDescription  
    }  
DESCRIPTION  
    "RAID Battery Missing."  
 ::= 84
```

everRun Enterprise User's Guide

```
controllerRAIDBatteryDegraded TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "RAID Battery Degraded."
    ::= 85
```

```
nodeNoRAIDDevices TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "No RAID Devices."
    ::= 86
```

```
controllerRAIDDiskOnNonRAIDController TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "RAID Disk On Non-RAID Controller."
    ::= 87
```

```
controllerMultipleLogicalDisks TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "Multiple Logical Disks."
    ::= 88
```

```
diskInvalidRAIDLevel TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
```

```
DESCRIPTION
"Invalid RAID Level."
::= 89

controllerMultiDiskRAID0BootDevice TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"MultiDisk RAID-0 Boot Device."
::= 90

diskBootDiskTooLarge TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"Boot Disk Too Large."
::= 91

nodeFirmwareNotSupported TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"Firmware Not Supported."
::= 92

controllerRAIDCapacitorFailed TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"RAID Capacitor Failed."
::= 93

controllerRAIDCapacitorMissing TRAP-TYPE
```

everRun Enterprise User's Guide

```
ENTERPRISE    everRunTrapId
VARIABLES    {
                everRunTrapDescription
            }
DESCRIPTION
"RAID Capacitor Missing."
::= 94

controllerRAIDCapacitorDegraded TRAP-TYPE
ENTERPRISE    everRunTrapId
VARIABLES    {
                everRunTrapDescription
            }
DESCRIPTION
"RAID Capacitor Degraded."
::= 95

nodeBmcConnectivity TRAP-TYPE
ENTERPRISE    everRunTrapId
VARIABLES    {
                everRunTrapDescription
            }
DESCRIPTION
"BMC Connectivity."
::= 96

diskDegraded TRAP-TYPE
ENTERPRISE    everRunTrapId
VARIABLES    {
                everRunTrapDescription
            }
DESCRIPTION
"Logical Disk Is Degraded."
::= 97

networkMiswired TRAP-TYPE
ENTERPRISE    everRunTrapId
VARIABLES    {
                everRunTrapDescription
            }
DESCRIPTION
```

```
"A Shared Network is miswired."
 ::= 98

networkNoBizPeerPort TRAP-TYPE
    ENTERPRISE  everRunTrapId
    VARIABLES   {
                everRunTrapDescription
                }
    DESCRIPTION
    "network_noBizPeerPort."
 ::= 99

unitNoFastSyncNetworkAvailable TRAP-TYPE
    ENTERPRISE  everRunTrapId
    VARIABLES   {
                everRunTrapDescription
                }
    DESCRIPTION
    "All DRDB sync networks are broken."
 ::= 100

networkCannotAutoCreateSharedNetwork TRAP-TYPE
    ENTERPRISE  everRunTrapId
    VARIABLES   {
                everRunTrapDescription
                }
    DESCRIPTION
    "network_cannot_auto_create_sharedNetwork."
 ::= 101

networkSlowSync TRAP-TYPE
    ENTERPRISE  everRunTrapId
    VARIABLES   {
                everRunTrapDescription
                }
    DESCRIPTION
    "network_slowSync."
 ::= 102

nodeIncorrectVNICSetting TRAP-TYPE
    ENTERPRISE  everRunTrapId
```

everRun Enterprise User's Guide

```
VARIABLES    {
                everRunTrapDescription
            }
DESCRIPTION
"node_incorrectVNICSetting."
::= 103

nodeIncorrectIMMSetting TRAP-TYPE
ENTERPRISE    everRunTrapId
VARIABLES    {
                everRunTrapDescription
            }
DESCRIPTION
"node_incorrectIMMSetting."
::= 104

unitLicenseSubscriptionExpired TRAP-TYPE
ENTERPRISE    everRunTrapId
VARIABLES    {
                everRunTrapDescription
            }
DESCRIPTION
"unit_licenseSubscriptionExpired."
::= 105

unitLicenseServiceExpired TRAP-TYPE
ENTERPRISE    everRunTrapId
VARIABLES    {
                everRunTrapDescription
            }
DESCRIPTION
"unit_licenseServiceExpired."
::= 106

unitLicenseAlasPollingFailed TRAP-TYPE
ENTERPRISE    everRunTrapId
VARIABLES    {
                everRunTrapDescription
            }
DESCRIPTION
"unit_licenseAlasPollingFailed."
```

```
 ::= 107

unitLicenseInvalidated TRAP-TYPE
    ENTERPRISE    everRunTrapId
    VARIABLES     {
        everRunTrapDescription
    }
    DESCRIPTION
        "unit_licenseInvalidated."
 ::= 108
```

```
unitLicenseServiceExpiryUnknown TRAP-TYPE
    ENTERPRISE    everRunTrapId
    VARIABLES     {
        everRunTrapDescription
    }
    DESCRIPTION
        "unit_licenseServiceExpiryUnknown."
 ::= 109
```

```
vmCannotRunLostDataAccess TRAP-TYPE
    ENTERPRISE    everRunTrapId
    VARIABLES     {
        everRunTrapDescription
    }
    DESCRIPTION
        "vm_cannot_run_no_data_access."
 ::= 110
```

```
unitLicenseUnsupportedPlatform TRAP-TYPE
    ENTERPRISE    everRunTrapId
    VARIABLES     {
        everRunTrapDescription
    }
    DESCRIPTION
        "unit_licenseUnsupportedPlatform."
 ::= 111
```

```
nodeUserPowerCycleRequired TRAP-TYPE
    ENTERPRISE    everRunTrapId
    VARIABLES     {
```

everRun Enterprise User's Guide

```
        everRunTrapDescription
    }
DESCRIPTION
"node_userPowerCycleRequired."
::= 112

nodeUserPowerOffRequired TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"node_userPowerOffRequired."
::= 113

nodeKernelDiagnosticPresent TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"node_kernelDiagnosticPresent."
::= 114

nodeReprovisionDom0NeedReboot TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"Reprovision Dom0 Need Reboot."
::= 115

nodeImsSingleLogicalDisk TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"IMS System Disk is Not Redundant."
::= 116
```

```
unitIsSyncing TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "unit_isSyncing."
    ::= 117
```

```
unitTestAlert TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "unit_testAlert."
    ::= 119
```

```
unitUnbalancedLoad TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "The Unit is not well balanced."
    ::= 120
```

```
unitNoAltSyncNetworks TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "unit_noAltSyncNetworks."
    ::= 121
```

```
unitNeedRepairStorage TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
```

everRun Enterprise User's Guide

```
    }
    DESCRIPTION
    "unit_needRepairStorage."
    ::= 122

localvmBlacklist TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "VMBlacklist."
    ::= 123

unitTooFew10GSyncLinks TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "unit_tooFew10GSyncLinks."
    ::= 124

unitTooFew1GSyncLinks TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "unit_tooFew1GSyncLinks."
    ::= 125

diskForeign TRAP-TYPE
    ENTERPRISE everRunTrapId
    VARIABLES {
        everRunTrapDescription
    }
    DESCRIPTION
    "disk_foreign."
    ::= 126
```

```
nodeNeedAddStorage TRAP-TYPE
    ENTERPRISE    everRunTrapId
    VARIABLES     {
        everRunTrapDescription
    }
    DESCRIPTION
    "node_needAddStorage."
    ::= 127
```

```
nodeCannotUpgrade TRAP-TYPE
    ENTERPRISE    everRunTrapId
    VARIABLES     {
        everRunTrapDescription
    }
    DESCRIPTION
    "node_cannotUpgrade."
    ::= 128
```

```
nodeCannotWorkOn TRAP-TYPE
    ENTERPRISE    everRunTrapId
    VARIABLES     {
        everRunTrapDescription
    }
    DESCRIPTION
    "node_cannotWorkOn."
    ::= 129
```

```
nodeCannotWorkOff TRAP-TYPE
    ENTERPRISE    everRunTrapId
    VARIABLES     {
        everRunTrapDescription
    }
    DESCRIPTION
    "node_cannotWorkOff."
    ::= 130
```

```
unitP2vFailed TRAP-TYPE
    ENTERPRISE    everRunTrapId
    VARIABLES     {
        everRunTrapDescription
    }
```

everRun Enterprise User's Guide

```
DESCRIPTION
"unit_p2vFailed."
::= 131

nodeSingleSystemDisk TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"node_singleSystemDisk."
::= 132

diskHasBadBlocks TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"disk_hasBadBlocks."
::= 133

nodeVolumeFailed TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"node_volumeFailed."
::= 134

unitVolumeFailed TRAP-TYPE
ENTERPRISE everRunTrapId
VARIABLES {
    everRunTrapDescription
}
DESCRIPTION
"unit_volumeFailed."
::= 135
```

-- End-of MIB(everRun) Revision 1.
-- End-of MIB(everRun) Revision 1.

-- =====
-- =====

END

