

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 and ESXi are 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 User's Guide*

Product Release Number: everRun Release 7.4.3.2

Publication Date: Wednesday, January 16, 2019

Stratus Technologies, Inc.

5 Mill and Main Place, Suite 500

Maynard, Massachusetts 01754-2660

© 2017 Stratus Technologies Bermuda, Ltd. All rights reserved.

Table of Contents

Part 1: everRun User's Guide	1
Chapter 1: Introduction to everRun Systems	1
everRun Quick Start Guide	1
Assembling Required Items	2
Configuring Your RAID Controller	2
Cabling Your System	3
Obtaining the ISO Image	4
Installing the everRun Software	4
Logging On to the everRun Availability Console	6
Creating a Protected Virtual Machine	7
everRun System Overview	7
everRun System Description	8
Physical Machines and Virtual Machines	9
Administrative Operations	9
Alerts	9
Remote Support	10
Lights-Out Management	10
Third-party Management Tools	10
Modes of Operation	11
High Availability Operation	11
Fault Tolerant Operation	12
SplitSite Configurations	13
SplitSite and Quorum Service	13
Quorum Servers	14
everRun Storage Architecture	15
Logical Disks and Physical Disks	15
Storage Groups	16
Sizing Volume Containers	16
External Storage	17
Network Architecture	18
Network Architecture Overview	19
A-Link and Private Networks	19

Business and Management Networks	20
Network Segmentation Fault Detection and Remediation	21
System Usage Restrictions	21
QEMU	21
Accessing the Host Operating System	21
Chapter 2: Getting Started	23
Planning	23
System Requirements Overview	24
System Hardware	24
Supported Servers	24
RAM	24
Disk Space Requirements	24
Network	24
IP Addresses	25
Ports	25
System Software	25
Storage Requirements	26
Memory Requirements	27
General Network Requirements and Configurations	27
Requirements	27
Recommended Configurations	28
Business and Management Network Requirements	28
A-Link and Private Network Requirements	30
SplitSite Network Requirements	30
General Network Requirements	30
A-Link Network Requirements	31
Business Network Requirements	31
Management Network Requirements	31
everRun Availability Console Requirements	32
Compatible Internet Browsers	32
Quorum Servers Considerations	32
Power Requirements and Considerations	34
Software Installation	34
Site and System Preparation	34

Connecting Power	35
UPS (Optional)	35
Obtaining everRun Software	36
Obtaining the ISO Image	36
Final Step	37
Creating Bootable USB Media	37
Configuring Settings in the Firmware Setup Utility	40
Required Settings	40
Recommended Settings	41
Installing everRun Software	41
Connecting Ethernet Cables	42
Installation Options	43
Installing Software on the First PM	45
Mapping Your Keyboard	48
To configure your keyboard layout during installation:	49
To configure your keyboard layout after installation:	49
Recording the Management IP Address	50
Installing Software on the Second PM	50
Post-Installation Tasks	53
Obtaining System IP Information	53
Logging on to the everRun Availability Console for the First Time	54
Connecting Additional Networks	55
Chapter 3: Using the everRun Availability Console	57
The everRun Availability Console	58
Logging on to the everRun Availability Console	59
The Dashboard Page	60
Resolving Outstanding Alerts on the Dashboard	60
The System Page	61
Rebooting the System	61
Shutting Down the System	62
The Preferences Page	63
Specifying Owner Information	66
Managing the everRun Product License	67
Configuring IP Settings	71

Configuring Quorum Servers	73
Configuring Date and Time	75
Configuring System Resources	76
Configuring Active Directory	77
Configuring the Migration Policy	79
Configuring Secure Connections	79
Configuring the Host Inactivity Logout	80
Disabling and Enabling Snapshots	81
Managing IPtables	82
Managing Diagnostic Files	85
Creating a Diagnostic File	85
Uploading a Diagnostic File to Customer Support	86
Deleting a Diagnostic File	87
Configuring e-Alerts	87
Configuring SNMP Settings	89
Configuring Remote Support Settings	91
Configuring Internet Proxy Settings	93
The Alerts Page	94
The Audits Page	95
The Physical Machines Page	95
Physical Machine Actions	96
Physical Machine States and Activities	98
The Virtual Machines Page	99
Virtual Machine Actions	100
Virtual Machine States and Activities	102
The Snapshots Page	104
The Volumes Page	104
The Storage Groups Page	106
The Networks Page	107
Fixing a Network Connection	108
The Virtual CDs Page	108
The Upgrade Kits Page	109
The Users & Groups Page	109
Managing Local User Accounts	110

Managing Domain User Accounts	112
Chapter 4: Upgrading everRun Software	115
Upgrading a System Using an Upgrade Kit	115
Upgrading a System To Release 7.4.x with a DVD	117
Chapter 5: Managing Logical Disks	119
Logical Disk Management	119
Responding to a Failed Logical Disk	120
Activating a New Logical Disk	122
Creating a New Storage Group	123
Deleting a Storage Group	123
Assigning a Logical Disk to a Storage Group	124
Chapter 6: Managing Physical Machines	125
Maintenance Mode	125
Rebooting a Physical Machine	127
Shutting Down a Physical Machine	127
Load Balancing	129
Modes of Operation	129
Troubleshooting Physical Machines	129
Recovering a Failed Physical Machine	130
Recovering a Failed Physical Machine with External Storage	132
Resetting MTBF for a Failed Physical Machine	132
Chapter 7: Managing Virtual Machines	135
Planning Virtual Machine Resources	136
Planning Virtual Machine vCPUs	136
Planning Virtual Machine Memory	138
Planning Virtual Machine Storage	139
Planning Virtual Machine Networks	141
Creating and Migrating Virtual Machines	141
Creating a New Virtual Machine	143
Creating a New Windows Server 2003 Virtual Machine	147
Copying a Virtual Machine	149
Migrating a Physical Machine or Virtual Machine to an everRun 7.x System	151
Migrating From Avance or everRun MX Systems	160
Planning to Migrate from an everRun MX System	161

Platform Requirements	162
Planned Outage	162
Guest Operating System Support	162
Network Preparation	162
Storage Considerations	163
Quorum Support	163
Installing everRun	164
Migrating Virtual Machines	164
Converting an everRun MX System to an everRun 7.x System	164
Planning to Migrate from an Avance Unit	171
Platform Requirements	171
Planned Outage	171
Guest Operating System Support	172
Network Preparation	172
Storage Considerations	173
Installing everRun	173
Migrating Virtual Machines	173
Converting an Avance Unit to an everRun 7.x System	173
Migrating a Windows Server 2003 VM to an everRun 7.x System	179
Importing an OVF File from an everRun MX System	182
Importing an OVF File from an Avance System	190
Importing an OVF File from an everRun 7.x System	198
Replacing a Virtual Machine from an OVF File	203
Exporting a Virtual Machine from an everRun System	208
Managing Windows Drive Labels	214
Configuring Windows-based Virtual Machines	215
Updating the VirtIO Drivers (Windows-based VMs)	215
Creating and Initializing a Disk (Windows-based VMs)	218
Installing Applications (Windows-based VMs)	219
Installing the QEMU Guest Agent for Application-Consistent Snapshots (Windows-based VMs)	219
Configuring Linux-based Virtual Machines	221
Creating and Initializing a Disk (Linux-based VMs)	222
Installing Applications (Linux-based VMs)	223

Installing the QEMU Guest Agent for Application-Consistent Snapshots (Linux-based VMs)	223
Managing the Operation of a Virtual Machine	224
Starting a Virtual Machine	225
Shutting Down a Virtual Machine	225
Powering Off a Virtual Machine	226
Opening a Virtual Machine Console Session	227
Renaming a Virtual Machine	228
Removing a Virtual Machine	229
Managing Virtual Machine Resources	230
Reprovisioning Virtual Machine Resources	231
Creating a Volume in a Virtual Machine	234
Attaching a Volume to a Virtual Machine	235
Detaching a Volume from a Virtual Machine	237
Removing a Volume from a Virtual Machine	238
Renaming a Volume on the everRun System	240
Expanding a Volume Container on the everRun System	241
Expanding a Volume on the everRun System	242
Recovering Virtual Machine Resources	243
Managing Virtual CDs	244
Creating a Virtual CD	244
Inserting a Virtual CD	246
Ejecting a Virtual CD	246
Booting from a Virtual CD	247
Renaming a Virtual CD	248
Removing a Virtual CD	248
Managing Snapshots	249
Creating a Snapshot	250
Creating a Virtual Machine from a Snapshot	254
Exporting a Snapshot	256
Removing a Snapshot	262
Advanced Topics (Virtual Machines)	263
Assigning a Specific MAC Address to a Virtual Machine	263
Selecting a Preferred PM for a Virtual Machine	265
Changing the Protection Level for a Virtual Machine (HA or FT)	265

Configuring the Boot Sequence for Virtual Machines	266
Resetting MTBF for a Failed Virtual Machine	267
Locating a Dump File in a Virtual Machine	268
Chapter 8: Maintaining Physical Machines	269
Physical Machine Hardware Maintenance Restrictions	270
Adding or Replacing Hot-Swappable Components	270
Adding or Replacing Components That Are Not Hot-Swappable	271
Adding a New NIC	272
Replacing Physical Machines, Motherboards, NICs, or RAID Controllers	274
Replacing a Failed Physical Machine Connected to External Storage	277
Upgrading Both Physical Machines In a Running System	278
Chapter 9: Managing External Storage	279
Adding or Removing External Storage LUNs	279
Configuring External Storage in a Newly Installed Release 7.4.x System	280
Configuring Linux Multipathing	281
Upgrading a System with External Storage from Release 7.3.4 to 7.4.x	284
Removing External Storage Disks	285
Updating initrd for Multipath	285
Part 2: Supporting Documents	287
Chapter 10: everRun Release 7.4.3.2 Release Notes	288
Accessing Stratus Knowledge Base Articles	288
Bug Fixes	289
Fixed in everRun Release 7.4.3.2	289
Fixed in everRun Release 7.4.3.1	290
Fixed in everRun Release 7.4.3.0	290
Fixed in everRun Release 7.4.2.0	291
Fixed in everRun Release 7.4.1.2	293
Fixed in everRun Release 7.4.1.1	293
Fixed in everRun Release 7.4.1.0	293
Fixed in everRun Release 7.4.0.1	294
New Features and Enhancements	294
New in everRun Release 7.4.3.0	294
New in everRun Release 7.4.2.0	294
New in everRun Release 7.4.1.0	295

New in everRun Release 7.4.0.0	295
Features No Longer Supported in everRun Release 7.4.0.0 or Higher	297
Important Considerations	297
Upgrading From everRun Release 7.4.0.x	298
Upgrading From everRun Release 7.3.4.0 or Previous Releases of everRun	298
Using 4K Disks for Better Performance	300
Do Not Update CentOS Host OS Directly From CentOS	301
Optimizing Performance of A-Link Networks	301
Status of RAID Physical Disks Is Not Monitored	301
Other Important everRun Considerations	301
Known Issues	301
Avoid Removing a VM While Only One Node is Running	301
Invalid VM Names Prevent Some VM Operations from Executing	301
Windows 2008 (pre-R2) Guests May Crash	302
Alert Status For a Failed Volume May Not Be Displayed Correctly	302
everRun May Not Detect Reconnection of Management Network Cable	302
Information About PCI Devices May Not Be Displayed	302
State of External Storage Not Reflected in Node's State	302
Resizing Logical Disks on External Storage Is Not Supported	302
VMs May Not Boot If One Node is Removed From System	303
Removing Snapshots Temporarily Prevents Some VM Operations	303
Creation of Snapshots Results in Volume Conversion from RAW to QCOW2 Format	303
Moving an everRun System to a Different Subnet	304
Enabling Log Files for the Windows QEMU Guest Agent May Cause VM Timeouts When Taking Snapshots	304
Unsupported Network Adapter Card and Chip	304
Do Not Use the ifdown Command	304
Documentation Updates	304
Updates in everRun Release 7.4.2.0 Documentation	304
Updates in everRun Release 7.4.1.0 Documentation	305
Updates in everRun Release 7.4.0.0 Documentation	305
Getting Help	307
Chapter 11: everRun Command Line Interface Reference	308
AVCLI Command Overview	308

Prerequisites	309
Installing the Linux Client	309
Installing the Windows Client	310
Using AVCLI	311
Executing a Command	311
Using AVCLI Help	312
Listing All Commands	312
Displaying Help for a Specific Command	313
AVCLI Error Status	313
XML Encapsulated Errors	313
Error Checking	314
Asynchronous Command Delay	314
Output Formatting	315
User-Friendly Command Output	315
Program-Friendly XML Output	316
AVCLI Exceptions	319
AVCLI Command Descriptions	320
ad-disable	326
ad-enable	327
ad-info	328
ad-join	329
ad-remove	330
alert-delete	331
alert-info	332
audit-export	333
audit-info	334
callhome-disable	335
callhome-enable	336
callhome-info	337
datetime-config	338
diagnostic-create	341
diagnostic-delete	342
diagnostic-extract	343
diagnostic-fetch	344

diagnostic-info	345
dialin-disable	346
dialin-enable	347
dialin-info	348
disk-move-to-group	349
ealert-config	350
ealert-disable	352
ealert-enable	353
ealert-info	354
help	355
image-container-info	356
image-container-resize	359
kit-delete	360
kit-info	361
license-info	362
license-install	363
local-group-add	364
local-group-delete	365
local-group-edit	366
local-group-info	367
local-user-add	368
local-user-delete	370
local-user-edit	371
local-user-info	373
localvm-clear-mtbf	374
media-create	375
media-delete	376
media-eject	377
media-import	378
media-info	380
media-insert	381
network-change-mtu	382
network-change-role	383
network-info	384

node-add	386
node-cancel	387
node-config-prp	388
node-delete	389
node-delete-prp	390
node-info	391
node-reboot	392
node-recover	393
node-shutdown	394
node-workoff	395
node-workon	396
ntp-config	397
ntp-disable	398
ova-info	399
ovf-info	400
owner-config	401
owner-info	402
pm-clear-mtbf	403
proxy-config	404
proxy-disable	405
proxy-enable	406
proxy-info	407
snmp-config	408
snmp-disable	410
snmp-info	411
storage-group-info	412
storage-info	413
timezone-config	414
timezone-info	415
unit-avoid-bad-node	416
unit-change-ip	417
unit-configure	419
unit-eula-accept	420
unit-eula-reset	421

unit-info	422
unit-shutdown	423
unit-shutdown-cancel	424
unit-shutdown-state	425
unit-synced	426
vm-ax-disable	427
vm-ax-enable	428
vm-boot-attributes	429
vm-cd-boot	430
vm-copy	431
vm-create	435
vm-create-from-snapshot	440
vm-delete	442
vm-export	443
vm-import	444
vm-info	447
vm-network-disable	448
vm-network-enable	449
vm-poweroff	450
vm-poweron	451
vm-reprovision	452
vm-restore	456
vm-shutdown	459
vm-snapshot-create	460
vm-snapshot-create-disable	462
vm-snapshot-create-enable	463
vm-snapshot-delete	464
vm-snapshot-export	465
vm-snapshot-info	467
vm-unlock	468
vm-volume-disable	469
vm-volume-enable	470
volume-info	471
volume-resize	472

Chapter 12: System Reference Information	474
Compatible Guest Operating Systems	474
Physical Machine System Requirements	476
Important Physical Machine and Virtual Machine Considerations	478
Virtual Machine Recommendations and Limits	478
Recommended Number of CPU Cores	478
Virtual Machine Limits	479
Combined Virtual Machine Maximums	480
Important Considerations	480
Accessing Knowledge Base Articles	481
Chapter 13: SNMP	484
Obtaining System Information with snmptable	484

Part 1: everRun User's Guide

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

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

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

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

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

For planning and installation information, see:

- ["Getting Started" on page 23](#)

The following topics describe how to administer everRun systems.

- ["Using the everRun Availability Console" on page 57](#)
- ["Upgrading everRun Software" on page 115](#)
- ["Migrating From Avance or everRun MX Systems" on page 160](#)
- ["Managing Logical Disks" on page 119](#)
- ["Managing Physical Machines" on page 125](#)
- ["Managing Virtual Machines" on page 135](#)

- ["Maintaining Physical Machines" on page 269](#)
- ["Managing External Storage " on page 279](#)

1

Chapter 1: Introduction to everRun Systems

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

See the following topics for an introduction to everRun systems:

- ["everRun System Overview" on page 7](#)
- ["Modes of Operation" on page 11](#)
- ["everRun Storage Architecture" on page 15](#)
- ["Network Architecture" on page 18](#)
- ["System Usage Restrictions" on page 21](#)

everRun Quick Start Guide

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

An everRun 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](#)
- ["Obtaining the ISO Image" on page 4](#)

- "Installing the everRun Software" on page 4
- "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 **Downloads** 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 24.
- 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 Availability Console Requirements" on page 32.
- A monitor, keyboard, and cables, used only during the installation.
- The everRun license key you received from Stratus.
- The everRun ISO image available at the **Downloads** page at <http://www.stratus.com/go/support/everrun>. Click **everRun** (if it is not already displayed) and then select the appropriate version. Scroll down to **Product Downloads**, and under **Install**, click the link to the appropriate ISO Image.
- From your network administrator, the IPv4 address, netmask, default gateway address, and DNS address values for everRun and each PM.

Configuring Your RAID Controller

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

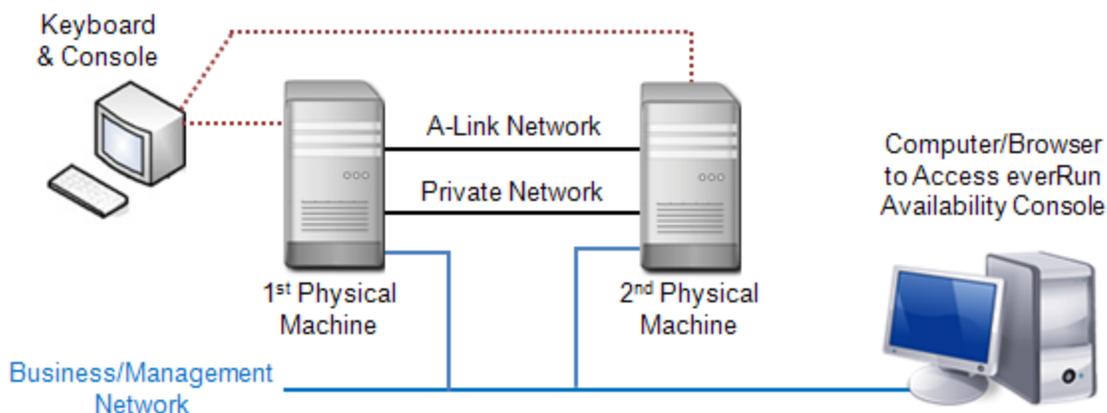
- If the system has a single logical disk, Stratus strongly recommends that you configure the RAID controller so that logical disks presented to the host are backed by redundant physical drives.
- Stratus strongly recommends that RAID controllers have a battery-backed write cache.
- You must configure 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 19](#).
- 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 34](#) 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.

Obtaining the ISO Image

Obtain the ISO image and create a bootable medium:

1. From any computer connected to the Internet, go to the **Downloads** page at <http://www.stratus.com/go/support/everrun>.
2. Click **everRun** (if it is not already displayed) and then select the appropriate version. Scroll down to **Product Downloads**, and under **Install**, click the link to the appropriate ISO image (everRun_install-7.x.x.x-xxx.iso). Save the ISO image.
3. Create a bootable medium using the ISO image:
 - To burn the ISO image to a DVD, use a commonly available application. For example, if you have the Roxio application installed, right click the ISO file and select the option to burn a DVD.
 - To create a bootable USB medium, see "[Creating Bootable USB Media](#)" on page 37.

Occasionally an ISO file can be corrupted during the download process. You can choose to verify the installation medium during the software installation.

For additional information, see "[Obtaining everRun Software](#)" on page 36.

Installing the everRun Software

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

1. Install the everRun software on the first PM:
 - a. Power on the first PM, and then insert the bootable medium or mount the ISO image.
 - b. As the PM powers on, configure the following settings in the firmware (BIOS or UEFI) setup utility:

- Set the first boot device to the Optical Drive.
- Enable Virtualization Technology.
- Enable Execute-Disable Bit Capability.



Note: If you need to configure your keyboard for a different layout, see "[Mapping Your Keyboard](#)" on page 48.

- c. From the installation software's Welcome screen, use the arrow keys to select an installation option. Stratus recommends that, for an initial installation, you select **Verify medium and Install everRun**, which verifies the installation medium before installing the software (verification adds about five minutes to the installation process). If you do not want to verify the medium, select **Install everRun, Create a new system**. After you select an option, press **Enter**.
- d. If any disk contains previously installed data, the following message appears and the system reboots (if no disk contains previously installed data, installation continues with the next step):

Rebooting because disks XXX were erased.

When the reboot is complete, the boot menu reappears and you must, again, select an installation option (Step 1,c, above).

- e. From the **Select interface for private Physical Machine connection** screen, select an embedded port (for example, **em1**, which is the default), and press **F12**.
- f. From the **Select interface for managing the system (ibiz0)** screen, select an embedded port (for example, **em2**, which is the default), and press **F12**.
- g. 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 50.

4. The **INITIAL CONFIGURATION** page appears. Under **NOTIFICATIONS**, the box for **Enable Support Notifications** is checked, by default. If you do not want the everRun 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 91](#)).
5. Under **SYSTEM IP**, for IP Address, enter the address you obtained from your network administrator. 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**.

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 244](#) in the online help.

Next, create a new virtual machine (VM) and install a guest operating system on your everRun 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 143](#) 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 53](#) in the online help.

everRun System Overview

An everRun 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 System Description" on page 8](#)
- ["Physical Machines and Virtual Machines" on page 9](#)
- ["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 System Description

everRun 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 systems. See ["Physical Machine System Requirements" on page 476](#) and ["System Requirements Overview" on page 24](#) 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 24](#)

["Compatible Guest Operating Systems" on page 474](#)

["Network Architecture Overview" on page 19](#)

Physical Machines and Virtual Machines

An everRun 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 57](#)

["The Physical Machines Page" on page 95](#)

["The Virtual Machines Page" on page 99](#)

Administrative Operations

You can perform many administrative operations on the everRun 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 58](#).

Alerts

everRun 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 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 91 for details.
- **Proxy Configuration**—Enables you to configure a proxy server for access to the Internet. See "[Configuring Internet Proxy Settings](#)" on page 93 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 systems fully support lights-out management on vendor servers.

Third-party Management Tools

You can install third-party management tools on everRun 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.
- 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 21](#).

Related Topics

["Getting Started" on page 23](#)

["System Reference Information" on page 474](#)

Modes of Operation

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

- ["High Availability Operation" on page 11](#)
- ["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 software provides two user-defined availability levels for VMs: High Availability (HA) and Fault Tolerant (FT).

In HA operation, everRun 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 99](#)

["Using the everRun Availability Console" on page 57](#)

Fault Tolerant Operation

everRun 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 99](#)

["Using the everRun Availability Console" on page 57](#)

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 30](#) 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 32](#)). 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 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 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 PMs and the quorum-service computers. For this to work properly, you must configure each everRun 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 19](#)

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 environment. Stratus strongly recommends using quorum servers, especially for SplitSite operation. You can configure an everRun 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 32](#)

["Configuring Quorum Servers" on page 73](#)

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

everRun Storage Architecture

The RAID controllers in an everRun system create logical disks from the system's physical disks. The logical disks are collected into storage groups. Logical disks contain everRun system volumes and virtual machine (VM) volumes. Each volume resides in an associated container that includes the current data on that volume as well as any volume snapshots.

The two physical machines (PM) in an everRun system can have different storage capacities, but only the smaller capacity is available to the system. For example, if one PM has 1 TB of storage in a storage group and the other has 2 TB of storage in that same storage group, only 1 TB is available to the everRun system for that storage group.

everRun systems support internal and external disks.

For more information about everRun storage, see the following topics:

- ["Logical Disks and Physical Disks" on page 15](#)
- ["Storage Groups" on page 16](#)
- ["External Storage" on page 17](#)
- ["Managing External Storage " on page 279](#)
- ["Sizing Volume Containers" on page 16](#)
- ["Storage Requirements" on page 26](#)

Logical Disks and Physical Disks

In an everRun 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 119](#).

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

["Storage Requirements" on page 26](#)

["The everRun Availability Console" on page 58](#)

["Using the everRun Availability Console" on page 57](#)

Storage Groups

In an everRun system, a storage group is a collection of logical disks. Multiple storage groups are supported. At install time, everRun software creates the **Initial Storage Group** which contains only the logical disk on which the software is being installed. After installation you can add other disks to any existing storage groups. If a logical disk is empty, you can move it to a different storage group.

With multiple storage groups, you can match application performance requirements to disk capabilities. You can group slower disks in one storage group and higher performance disks in another storage group. You can then assign the volumes of VMs that run more demanding applications to a storage group with higher performance disks.

You can view information about storage group from the **Storage Groups** page of the everRun Availability Console. For information, see ["The Storage Groups Page" on page 106](#).

Related Topics

["Storage Requirements" on page 26](#)

["Creating a New Storage Group" on page 123](#)

["The everRun Availability Console" on page 58](#)

["Using the everRun Availability Console" on page 57](#)

Sizing Volume Containers

A *volume container* is storage space that holds a volume and VM snapshot data associated with that volume.

You can specify the size of the volume container when you create a VM. As snapshot data accumulates, you may need to increase the size of the volume container. You can expand the volume container, but you cannot decrease its size.

The following factors affect the size of a volume container:

- The volume size
- If snapshots are being taken:
 - The number of snapshots retained
 - How much data changes between snapshots



Note: The amount of data that changes between snapshots varies for different applications and can have a large impact on what size the volume container should be. In order to properly size a volume container, you must consider the amount of data that your application changes between snapshots.

If you do not take snapshots, the size of the volume container can be the volume size.

If you take snapshots, the size of the volume container depends largely on the amount of data that is written to the volume between snapshots:

- For a VM created with a separate boot disk, or for applications that write relatively small amounts of data between snapshots, a reasonable volume-container size is 2.6 times larger than the volume size.
- For applications that write moderate amounts of data between snapshots, a reasonable volume-container size is approximately 3.5 times larger than the volume size.
- For applications that write larger amounts of data between snapshots, the volume-container size must be more than 3.5 times larger than the volume size.

The following is a general formula to calculate the *approximate* volume container size :

$$VolContSize = 2 * VolSize + [(\# SnapshotsRetained + 1) * SnapshotSize]$$

Related Topics

["Expanding a Volume Container on the everRun System" on page 241](#)

["image-container-resize" on page 359](#)

External Storage

everRun systems support Fibre Channel, iSCSI, and Linux device mapper multipath external storage.

everRun software handles external storage in the same way it handles internal storage. It mirrors VM data between external storage logical disks attached to each node as it does for internal storage. This is known as *everRun replicated* storage.

everRun software has no external storage configuration feature. You must manage the installation and configuration of your particular external storage infrastructure. This may involve installing HBA cards, drivers, and tools required by your external storage system. Follow the instructions provided by your external

storage system vendor. In addition you must configure the everRun system to use Linux multipathing in order for the system to support external storage systems (see ["Configuring Linux Multipathing" on page 281](#)).



Caution: everRun boot partitions and other everRun system data must reside on internal storage; they are **not** supported on external storage.



Note: The system administrator must have root access to the everRun host OS on each physical machine in order to install and/or configure external storage.

For Fibre Channel external storage, this typically involves:

1. Installing and configuring HBAs.
2. Obtaining information (WWNs, port addresses, etc.) to be used in masking and zoning operations supported by your SAN fabric infrastructure.
3. Provisioning LUNs on the external storage system.
4. Managing access to logical disks by configuring SAN zoning and LUN masking.

For iSCSI external storage, this typically involves:

1. Configuring your iSCSI initiators.
2. Provisioning LUNs on the external storage system.
3. Managing access to logical disks by configuring LUN masking.

After installing and configuring your external storage infrastructure you must configure the everRun system host OS to recognize and use the external storage. See ["Adding or Removing External Storage LUNs" on page 279](#) for more information.

Related Topics

["Managing External Storage" on page 279](#)

Network Architecture

See the following topics for information about everRun network architecture.

- ["Network Architecture Overview" on page 19](#)
- ["A-Link and Private Networks" on page 19](#)

- ["Business and Management Networks" on page 20](#)
- ["Network Segmentation Fault Detection and Remediation" on page 21](#)

Network Architecture Overview

Ethernet networks provide pathways for communications between two physical machines (PMs) in an everRun 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 19](#).
- *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 20](#).

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

An everRun system also provides a network segmentation detection mechanism. For information, see ["Network Segmentation Fault Detection and Remediation" on page 21](#).

A-Link and Private Networks

Every everRun 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 `priv0`, 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 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 55](#).

Related Topics

["Business and Management Networks" on page 20](#)

["A-Link and Private Network Requirements" on page 30](#)

["Network Architecture Overview" on page 19](#)

["Fixing a Network Connection" on page 108](#)

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 PM has a single management network, referred to as *ibiz0*.

The everRun 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 55](#).

Related Topics

["A-Link and Private Networks" on page 19](#)

["Business and Management Network Requirements" on page 28](#)

["Network Architecture Overview" on page 19](#)

["Fixing a Network Connection" on page 108](#)

Network Segmentation Fault Detection and Remediation

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

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

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

Related Topics

["Network Architecture Overview" on page 19](#)

System Usage Restrictions

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

- ["QEMU" on page 21](#)
- ["Accessing the Host Operating System" on page 21](#)

QEMU

Stratus everRun 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 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 50](#).



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

The default 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 planning, installation and post-installation tasks:

- ["Planning" on page 23](#)
- ["Software Installation" on page 34](#)
- ["Post-Installation Tasks" on page 53](#)

Planning

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

- ["System Requirements Overview" on page 24](#)
- ["Storage Requirements" on page 26](#)
- ["Memory Requirements" on page 27](#)
- ["General Network Requirements and Configurations" on page 27](#)
- ["Business and Management Network Requirements" on page 28](#)
- ["A-Link and Private Network Requirements" on page 30](#)
- ["SplitSite Network Requirements" on page 30](#)
- ["everRun Availability Console Requirements" on page 32](#)
- ["Compatible Internet Browsers" on page 32](#)
- ["Quorum Servers Considerations" on page 32](#)
- ["Power Requirements and Considerations" on page 34](#)

After you have planned the system configuration, continue with ["Software Installation" on page 34](#).

System Requirements Overview

An everRun 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 ["System Hardware " on page 24](#) requirements are described below. See ["System Software" on page 25](#) for software requirements.

System Hardware

Supported Servers

Stratus everRun software will run on any systems listed on the [Red Hat® Linux Hardware Catalog](#) that support RHEL 7.x and any of the supported processors listed in ["Physical Machine System Requirements" on page 476](#).

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 firmware (BIOS or UEFI) setup utility.

RAM

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

Disk Space Requirements

Internal disks and external Fibre Channel and iSCSI disks are supported. A minimum of two drives per physical machine is required.

477 MB is required on each internal logical disk for the host CentOS operating system. In addition, 22 GB is required on two of the internal logical disks for everRun system data including logs. Only internal disks may be boot disks. The amount of disk space required for a VM's boot volume varies depending on the operating system being used. Additional storage is needed for applications and data on each VM, as well as VM snapshots.

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.

SplitSite configurations have different network requirements. For information, see ["SplitSite Network Requirements" on page 30](#).

See ["Network Architecture Overview" on page 19](#), ["A-Link and Private Networks" on page 19](#), and ["Business and Management Networks" on page 20](#) for more information.

IP Addresses

Each everRun 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 53](#) for more information.

Ports

everRun systems use port 443 in the local firewall for HTTPS communications, port 22 for ssh, and 5900-5999 for each active VNC associated with each VM. Firewalls must allow traffic through the appropriate ports. Firewalls must permit everRun-protected VMs to contact quorum service computers using UDP port 4557. For additional information on TCP and UDP ports, access the Knowledge Base to search for the article *everRun TCP and UDP ports*. See ["Accessing Knowledge Base Articles" on page 481](#).

System Software

See ["Compatible Guest Operating Systems" on page 474](#).

Related Topics

["Physical Machine System Requirements" on page 476](#)

["Important Physical Machine and Virtual Machine Considerations" on page 478](#)

["Virtual Machine Recommendations and Limits" on page 478](#)

["Planning Virtual Machine Resources" on page 136](#)

["Configuring IP Settings" on page 71](#)

Storage Requirements

An everRun system has the following storage requirements and recommendations:

- Each physical machine must contain at least two physical disks.
- Stratus strongly recommends that your system use 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.

Disk drives support standard format, 512e format, and Advanced 4K Native format with sectors, as follows:

Format	Physical Sectors	Logical Sectors
Standard	512B	512B
512e	4KiB	512B
Advanced 4K Native	4KiB	4KiB

Stratus recommends using disks with the 4K sector size for better performance. When using 4K storage, observe the following restrictions:

- Each storage group must contain logical disks of the same sector size. For example, you cannot add a 4K logical disk to a 512B storage group.
- The sector size of a storage group is automatically defined by the sector size of the first logical disk that you add to the group. After the sector size of the storage group is set, you cannot change it. You would need to create a new storage group with logical disks of the new sector size.
- Because the sector size of a storage group affects the sector size of your VM volumes, it is important to plan your storage groups carefully:

- A storage group with the 512B sector size supports only 512B VM volumes.
- A storage group with the 4K sector size supports 4K or 512B VM volumes.

Note that the boot volume for each VM must be 512B, regardless of the storage group's sector size. Only data volumes can use the 4K sector size. Ensure that your guest operating systems support 4K volumes before creating or attaching them.

When planning your system configuration, confirm that your storage configuration meets these requirements and then return to ["Site and System Preparation" on page 34](#).

Related Topics

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

["Managing Logical Disks" on page 119](#)

["The Storage Groups Page" on page 106](#)

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 (the least memory of either PM).

Related Topics

["Planning Virtual Machine Memory" on page 138](#)

General Network Requirements and Configurations

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

Requirements

Before you install everRun 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 30](#)
- ["Business and Management Network Requirements" on page 28](#)
- ["SplitSite Network Requirements" on page 30](#)

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:

- The networks use IPv6 link-local addressing.
- The speed of business or management networks should be less than or equal to the speed of A-Link networks.
- The networks support an MTU value of up to 9000.
- The networks do not support 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.

Related Topics

["Network Architecture Overview" on page 19](#)

["Business and Management Networks" on page 20](#)

["network-info" on page 384](#) and ["network-change-mtu" on page 382](#)

A-Link and Private Network Requirements

A-Link and private networks have the following requirements:

- The networks use IPv6 link-local addressing.
- All A-Link and private networks on one PM of 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.
- Each PM can have one to eight A-Link networks; however, a minimum of two is recommended.
- The networks use 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 19](#)

SplitSite Network Requirements

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

- ["A-Link Network Requirements" on page 31](#)
- ["SplitSite Network Requirements" on page 30](#)
- ["Business Network Requirements" on page 31](#)
- ["Management Network Requirements" on page 31](#)

General Network Requirements

All network components in a SplitSite configuration must have more than 155 Mbps minimum capacity, end-to-end.

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 maximum 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 maximum 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.
- When fault-tolerant SMP is used, the A-Link networks must be a minimum of 1Gbps.

Business Network Requirements

Business networks in a SplitSite configuration require the following:

- 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 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 19](#)

everRun 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 32](#).

For more information, see ["Using the everRun Availability Console" on page 57](#).

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 systems.

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

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 Server 2012, Windows Server 2008, Windows Server 2003, Windows Vista, Windows 7, or Windows 8; always powered on

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

- 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:

1. Open the **Downloads** page at <http://www.stratus.com/go/support/everrun>: scroll down to the **Software Downloads** section and then, for **Select Product & Version**, select **everRun** and the appropriate version number.
2. Scroll down to the **Drivers and Tools** section and then click **Quorum Service** to download the quorum server software installer file to the quorum server.
3. On the quorum server, double click the installer file.



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 30 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 9](#)

["Configuring Quorum Servers" on page 73](#)

["The everRun Availability Console" on page 58](#)

["Using the everRun Availability Console" on page 57](#)

Power Requirements and Considerations

To ensure maximum availability, Stratus strongly recommends that everRun'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 35](#) 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 software for the first time:

1. Prepare your site and system for the installation. See ["Site and System Preparation" on page 34](#).
2. Connect power to your system. See ["Connecting Power" on page 35](#).
3. Install the everRun software. See ["Installing everRun Software" on page 41](#).

When the software installation is complete, see ["Post-Installation Tasks" on page 53](#).

Related Topics

["Upgrading everRun Software" on page 115](#)

Site and System Preparation

Before you install everRun 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 24](#).
- The storage configuration meets all of the requirements described in ["Storage Requirements" on page 26](#).
- 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 Availability Console Requirements" on page 32](#).
- Determine the best configuration for your network. See ["General Network Requirements and Configurations" on page 27](#).
- Use either an internal DVD drive or bootable USB media (see ["Creating Bootable USB Media" on page 37](#)) for the installation.

After confirming that your site and system meet the preceding requirements, return to ["Software Installation" on page 34](#).

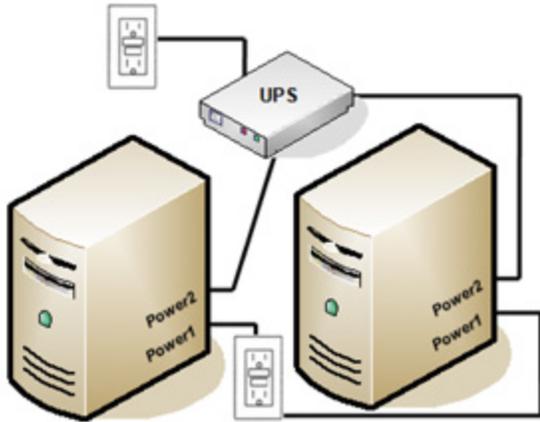
Connecting Power

To connect power, configure your everRun server with redundant power supplies connected to separate sources. After connecting power, return to ["Software Installation" on page 34](#).

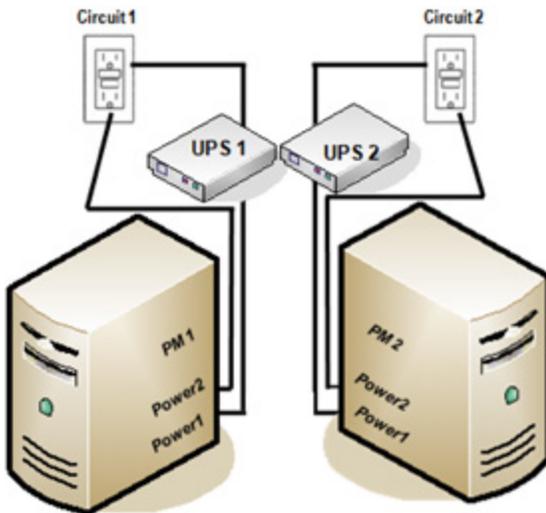
UPS (Optional)

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 34](#)

Obtaining everRun Software

Stratus provides the everRun software as an ISO image. You can boot from it directly, or you can create bootable media.

Obtaining the ISO Image

1. From any computer connected to the Internet, go to the **Downloads** page at <http://www.stratus.com/go/support/everrun>.

2. To download the everRun software ISO image (everRun_install-7.x.x.x-xxx.iso), under **Product Downloads**, click **everRun 7.x.x.x ISO Image**. Save the ISO image.



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

Final Step

After you have obtained the ISO image, do one of the following:

- Create bootable media. Burn the ISO image to a DVD using a commonly available application, or create a bootable USB medium (see ["Creating Bootable USB Media" on page 37](#)). Then, perform the next step in ["Installing everRun Software" on page 41](#).
- If you are not creating bootable media, perform the next step in ["Installing everRun Software" on page 41](#).

Occasionally an ISO file can be corrupted during the download process. You can choose to verify the installation medium during the software installation.

Creating Bootable USB Media

After you have saved the everRun installation software ISO image, you have the option of copying the image to a bootable USB medium, such as a thumbdrive. Follow the procedure below for your system.

Creating a bootable USB medium on a Linux-based system



Caution: This procedure will destroy any data on the USB medium.

1. Log in to the system as **root**.
2. Insert a USB medium such as a thumbdrive into the system. Determine the name of the thumbdrive.

One method of determining the name of the thumbdrive is to execute the **dmesg** command to display a log of all recent events, as in this example:

```
# dmesg | tail
```

The end of the log includes messages related to the recently inserted thumbdrive, as in this example:

```
sd 6:0:0:0: [sdc] Attached SCSI removable disk
```

Make note of the thumbdrive name in the messages (**sdc**, in the example).

3. Check if the system mounted the thumbdrive automatically.

One method of determining if the system mounted the thumbdrive automatically is to execute the **findmnt *thumbdrive_name*** command, as in this example:

```
# findmnt | grep sdc
```

If the command displays no output, the thumbdrive was not mounted and you can continue with the next step. If the command displays output, the system mounted the thumbdrive automatically and you must unmount it. Note the **TARGET** in the command output, as in this example:

```
TARGET SOURCE FSTYPE OPTIONS
/media/MY-DATA /dev/sdc1 vfat
```

Then, issue the command **umount *TARGET***, as in this example

```
# umount /media/MY-DATA
```

4. Write the installation software ISO image directly to the thumbdrive.

One method of writing the image is to execute the **dd** command in the format **dd if=*path_to_image* iso of=/dev/*sdx* bs=*blocksize***, where ***path_to_image*** is the full path to the ISO image file that you saved, ***sdx*** is the thumbdrive device name, and ***blocksize*** is an amount that ensures a timely writing process. The following command is an example:

```
# dd if=Downloads/everRun_install-7.4.0.0-97.iso
of=/dev/sdc bs=8K
```

Wait while the **dd** command completes processing. A prompt appears when the command is complete.

5. Log out and remove the thumbdrive. The thumbdrive is ready to be used as a boot device.

Creating a bootable USB medium on a Windows-based system



Caution: This procedure will destroy any data on the USB medium.



Note: Copying an ISO image to a USB medium using a file manager such as Windows Explorer or a similar tool does not create a bootable device.

Many utilities write an ISO image to a USB medium such as a thumbdrive on Windows-based systems. The following procedure uses the utility **Rufus**, which is available at <http://rufus.akeo.ie/>.

1. Download and save the everRun software ISO image (see "[Obtaining everRun Software](#)" on [page 36](#)) to a Windows-based system, if you have not already done so.
2. Ensure the integrity of the ISO image using a tool for verifying MD5 checksums on Windows-based systems. For example, use the **Microsoft File Checksum Integrity Verifier** tool, which is available at <https://www.microsoft.com/en-us/download/details.aspx?id=11533>.
3. Download and install the **Rufus** utility at <http://rufus.akeo.ie/>. Be sure to click the **Download** link that is about halfway down the web page (below **Last updated**); do not click advertisements, including **Download** links that appear in advertisements.
4. Insert a USB medium such as a thumbdrive into the system
5. Run the **Rufus** utility, selecting the following options:

Option	Value(s)
Partition scheme and target system type	MBR partition scheme for BIOS and UEFI
File system	FAT32
Cluster size	4096 bytes
Format Option	Create a bootable disk using ISO Image (browse to the xxx.iso image) Create extended label and icon files

6. Click **Start** after selecting the options.
7. On the menu that appears, select **Write in DD image mode**.
8. Click **OK** to write to the USB device.

When the utility has completed writing the USB stick, **READY** appears in the horizontal box near the bottom of the utility interface. You can remove the USB device and use it to install everRun.

When the USB device is ready to be used to install everRun software, perform the next step in ["Installing everRun Software" on page 41](#).

Related Topics

["Obtaining everRun Software" on page 36](#)

["Software Installation" on page 34](#)

Configuring Settings in the Firmware Setup Utility

Before you install the software, you must modify settings in the firmware (BIOS or UEFI) setup utility. You can also modify some optional, though recommended, settings.



Note: A system with UEFI firmware always boots from the original software boot disk. If the boot disk fails, you need to perform a recover node operation (see ["Recovering a Failed Physical Machine" on page 130](#)).

After you modify the settings, save them and perform the next step in the installation procedure (either ["Installing Software on the First PM" on page 45](#) or ["Installing Software on the Second PM" on page 50](#)).



Note: This topic provides general information about settings in the firmware setup utility. Because settings—including setting names—vary, refer to the manufacturer's documentation for specific instructions for modifying any setting.

Required Settings

The following settings are required.

First Boot Device	Controls which device boots the operating system. Set the first boot device to the Optical Drive.
Virtualization 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 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 Disable, as the everRun system may be able to provide more information once the server is running.

Installing everRun 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 software. See ["Obtaining everRun Software" on page 36](#)
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 34](#)).
 - b. Connect Ethernet cables for the networks you are configuring. See ["Connecting Ethernet Cables" on page 42](#).
3. Perform the installation on the first PM. See .
4. After you have finished installing the software on the first PM, perform the installation on the second PM. See .

The installation is now complete. To perform the required post-installation configuration steps, see ["Post-Installation Tasks" on page 53](#).

Connecting Ethernet Cables

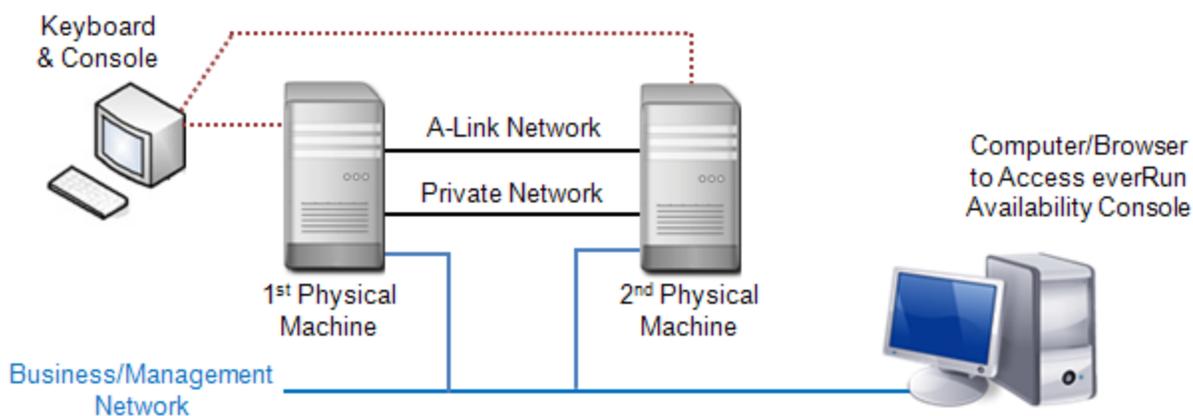
Before you install everRun software for the first time, you need to connect Ethernet cables for your networks.



Note: To install additional networks after you have completed the software installation, see ["Connecting Additional Networks" on page 55](#).

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 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 30

- 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 Software"](#) on page 41.

Related Topics

["A-Link and Private Network Requirements"](#) on page 30

["Business and Management Network Requirements"](#) on page 28

["everRun Availability Console Requirements"](#) on page 32

["Connecting Additional Networks"](#) on page 55

Installation Options

When you begin the installation, 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
Verify the installation medium and then perform the installation.	Verify medium and Install everRun	Verifies the installation medium first, then installs the CentOS and everRun software, and, finally, creates a new system. (Stratus recommends that you verify installation medium the first time you use it; note, though, that verification adds five or so minutes to the installation.) See .

Task	Option	Description
Verify the installation medium and then recover a physical machine.	Verify medium and Recover Physical Machine	Verifies the installation medium and then recovers a physical machine. See "Recovering a Failed Physical Machine" on page 130.
Verify the installation medium and then replace a physical machine.	Verify medium and Replace Physical Machine	Verifies the installation medium and then replaces a physical machine. See "Replacing Physical Machines, Motherboards, NICs, or RAID Controllers" on page 274.
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 .
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 130.
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 and "Replacing Physical Machines, Motherboards, NICs, or RAID Controllers" on page 274.

Task	Option	Description
Boot to rescue mode (UEFI firmware installation only)	Rescue the installed system	Boots to rescue mode.

Installing Software on the First PM

This topic describes how to perform an initial installation of the everRun 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 bootable medium or mount the ISO image.
2. As the system powers on, enter the firmware (BIOS or UEFI) setup utility, and configure the required and optional settings described in ["Configuring Settings in the Firmware Setup Utility"](#) on page 40.



Note: If you need to configure your keyboard for a different language, see ["Mapping Your Keyboard"](#) on page 48.

3. When the installation software loads, the **Welcome** screen appears and displays the options shown in ["Installation Options"](#) on page 43. 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

- Use the arrow keys to select an installation option. Stratus recommends that, for an initial installation, you select **Verify medium and Install everRun**, which verifies the installation medium before installing the software (verification adds about five minutes to the installation process). If you do not want to verify the medium, select **Install everRun, Create a new system**. After you select an option, press **Enter**.
- If any disk contains previously installed data, the following message appears and the system reboots (if no disk contains previously installed data, installation continues with the next step):

Rebooting because disks XXX were erased.

When the reboot is complete, the boot menu reappears and you must, again, select **Method 1** or **Method 2** (Step 3, above).

- 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.

- 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.

- v. 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.
- vi. 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

- i. Press the **Tab** key to bring up the command line.
- ii. 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**

iii. 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

iv. After entering the desired values on the command line, press **Enter**.

v. If any disk contains previously installed data, the following message appears and the system reboots (if no disk contains previously installed data, installation continues with the next step):

Rebooting because disks XXX were erased.

When the reboot is complete, the boot menu reappears and you must, again, select **Method 1** or **Method 2** (Step 3, above).

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 bootable medium, 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 50

5. Perform the next step in "[Installing everRun Software](#)" on page 41.

Mapping Your Keyboard

You can configure your keyboard for a different layout either during or after installation.

To configure your keyboard layout during installation:

Supported keyboard layouts include:

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

To configure your keyboard layout during installation:

1. As the first PM boots, select **Install**, **Recover**, or **Replace** from the boot menu.
2. On legacy BIOS systems, press **Tab** to access the kernel command line. On UEFI systems, press **e**.
3. Specify the `inst.keymap` kernel argument to configure the correct keyboard layout. The following example configures the Japanese keyboard layout:

```
inst.keymap=jp106
```
4. On legacy BIOS systems, press **Enter** to continue the boot sequence. On UEFI systems, press **Ctrl-x**.
5. Repeat the preceding steps on the second PM.

To configure your keyboard layout after installation:

1. Log in to the first PM as `root`.
2. From the command line, issue the `localectl` command to configure the correct keyboard layout. The following example configures the German keyboard layout:

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

3. Repeat the preceding steps on the second PM.

Related Topics

["Installing Software on the First PM" on page 45](#)

["Post-Installation Tasks" on page 53](#)

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

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 Software" on page 41](#) to see the next step.

Installing Software on the Second PM

This topic describes how to perform an initial installation of the everRun 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 bootable media or mount the ISO image.
2. As the system powers on, enter the firmware (BIOS or UEFI) setup utility, and configure the required and optional settings described in ["Configuring Settings in the Firmware Setup Utility" on page 40](#).
3. When the installation software loads, the **Welcome** screen appears and displays the options shown in ["Installation Options" on page 43](#). 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 45](#).
4. Use the arrow keys to select **Replace PM, Join system: Initialize data**, and press **Enter**. (If you verified the installation medium when installing the software on the first PM, you do not need to verify the installation medium now.)



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

5. If any disk contains previously installed data, the following message appears and the system reboots (if no disk contains previously installed data, installation continues with the next step):

Rebooting because disks XXX were erased.

When the reboot is complete, the boot menu reappears and you must continue with Step 3, above.

6. 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.

7. 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.

8. 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.
9. 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.

10. 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 bootable media, 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 50](#)
11. Perform the next step in ["Installing everRun Software" on page 41](#).

Post-Installation Tasks

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

- ["Obtaining System IP Information" on page 53](#)
- ["Logging on to the everRun Availability Console for the First Time" on page 54](#)
- Configuring Required System Preferences:
 - ["Configuring Date and Time" on page 75](#)
 - ["Configuring Remote Support Settings" on page 91](#)
 - ["Configuring Quorum Servers" on page 73](#)
 - ["Specifying Owner Information" on page 66](#)
- ["Configuring Active Directory" on page 77](#)
- ["Managing Local User Accounts" on page 110](#)
- ["Resolving Outstanding Alerts on the Dashboard" on page 60](#)
- ["Connecting Additional Networks" on page 55](#)

In addition, you must configure external storage if your system includes it. See ["Configuring External Storage in a Newly Installed Release 7.4.x System" on page 280](#).

Obtaining System IP Information

After you install the everRun 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](#)

54). 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 50](#)), which helps the network administrator determine system IP information. The system IP address must be a static IP address. Do not use a dynamic IP address.

Related Topics

["Software Installation" on page 34](#)

["Post-Installation Tasks" on page 53](#)

Logging on to the everRun Availability Console for the First Time

After completing the installation of the everRun software, log on to the everRun Availability Console to accept the end-user license agreement (EULA) and to manage the everRun 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 50](#).
- The system IP address—The network administrator provides this information. See ["Obtaining System IP Information" on page 53](#).
- The license **.KEY** file that you received from Stratus when you purchased the everRun 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 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 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 91).
5. Under **SYSTEM IP**, for IP Address, enter the address you obtained from your network administrator.

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 110.

Related Topics

["Software Installation" on page 34](#)

["The everRun Availability Console" on page 58](#)

["Using the everRun Availability Console" on page 57](#)

Connecting Additional Networks

The everRun 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 42](#)

["A-Link and Private Network Requirements" on page 30](#)

["Business and Management Network Requirements" on page 28](#)

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

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 system from a remote management computer. For an overview of the console, see "[The everRun Availability Console](#)" on page 58.

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

- "[The Dashboard Page](#)" on page 60
- "[The System Page](#)" on page 61
- "[The Preferences Page](#)" on page 63
- "[The Alerts Page](#)" on page 94
- "[The Audits Page](#)" on page 95
- "[The Physical Machines Page](#)" on page 95
- "[The Virtual Machines Page](#)" on page 99
- "[The Snapshots Page](#)" on page 104
- "[The Volumes Page](#)" on page 104
- "[The Storage Groups Page](#)" on page 106
- "[The Networks Page](#)" on page 107
- "[The Virtual CDs Page](#)" on page 108
- "[The Upgrade Kits Page](#)" on page 109
- "[The Users & Groups Page](#)" on page 109

The everRun Availability Console

The everRun Availability Console is a browser-based interface that provides management and monitoring of an everRun 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 Availability Console Requirements" on page 32](#).

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 60](#).
- View VM, CPU, memory, and storage statistics, and reboot or shutdown the system from the System page. See ["The System Page" on page 61](#).
- Set preferences for the system, diagnostics, notifications (e-Alerts and SNMP configuration), and remote support (notification and access); and access administrative tools that enable you to set a migration policy, create a secure connection, configure snapshots, and set other functionality. System preferences include owner information and configuration values for IP address, quorum services, date and time, etc. See ["The Preferences Page" on page 63](#).
- View alerts and audit logs. See ["The Alerts Page" on page 94](#) and ["The Audits Page" on page 95](#).
- Monitor, manage, and maintain resources:
 - PM status, storage, disks, network, and sensors: see ["The Physical Machines Page" on page 95](#).
 - VM status and management tasks such as creating, importing/restoring, managing, and maintaining VMs: see ["The Virtual Machines Page" on page 99](#).
 - Snapshot status and management tasks such as exporting and deleting snapshots: see ["The Snapshots Page" on page 104](#).
 - Volumes, including their state, size, and storage group: see ["The Volumes Page" on page 104](#).
 - Storage groups, including name, size used, size, and number of volumes: see ["The Storage Groups Page" on page 106](#).

- Networks, including state, physical interface, speed, MAC address, and network bandwidth: see ["The Networks Page" on page 107](#).
- Virtual CDs, including their state, name, size, and storage group: see ["The Virtual CDs Page" on page 108](#).
- Monitor and manage upgrade kits, and users and groups in the LIBRARY. See ["The Upgrade Kits Page" on page 109](#) and ["The Users & Groups Page" on page 109](#).

Related Topics

["Logging on to the everRun Availability Console for the First Time" on page 54](#)

["Logging on to the everRun Availability Console" on page 59](#)

["Using the everRun Availability Console" on page 57](#)

Logging on to the everRun Availability Console

Log on to the everRun Availability Console to manage the everRun 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 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 54](#)

["The everRun Availability Console" on page 58](#)

["Using the everRun Availability Console" on page 57](#)

The Dashboard Page

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

To display additional information about outstanding alerts, click an alert symbol (for example, ) in the everRun system diagram or click an entry in the list of alerts below the system diagram. Alert lists may appear in tabs such as **All**, **System**, or **Ignored**, which may appear below the system diagram, depending on the alerts. The alert information includes:

- The component associated with the issue (for example, the 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 60](#)).

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 95](#)

["The System Page" on page 61](#)

["The Virtual Machines Page" on page 99](#)

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 60](#)

The System Page

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

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

- ["Rebooting the System" on page 61](#)
- ["Shutting Down the System" on page 62](#)

You perform many other administrative tasks on the everRun system using the everRun Availability Console. For information, see ["The everRun Availability Console" on page 58](#).

To manage everRun system resources, see ["Configuring System Resources" on page 76](#).

Related Topics

["Using the everRun Availability Console" on page 57](#)

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. A message appears, asking you to confirm the reboot. Click **Yes** to continue.

Rebooting can take up to 15 minutes. You can observe the process in the **Dashboard** and the masthead of the everRun Availability Console. The system's PMs sequentially enter and then exit maintenance mode (for information on maintenance mode, see "[Maintenance Mode](#)" on page 125).

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 58](#)

["The System Page" on page 61](#)

["Using the everRun Availability Console" on page 57](#)

Shutting Down the System

Use the everRun Availability Console to shut down the everRun 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. A warning appears: *It will shut down the entire system and stop one or more VMs!* Click **Yes** to shutdown or **No** to cancel the shutdown. After clicking **Yes**, a second warning appears, asking you to confirm the shutdown. Click **Yes** (again) to shutdown or **No** to cancel the shutdown.

You can observe some of the shutdown process in the **Dashboard** and the masthead of the everRun Availability Console as the system's PMs sequentially enter maintenance mode (for information on maintenance mode, see ["Maintenance Mode" on page 125](#)). 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

["Managing the Operation of a Virtual Machine" on page 224](#)

["The everRun Availability Console" on page 58](#)

["The System Page" on page 61](#)

["Using the everRun Availability Console" on page 57](#)

The Preferences Page

The **Preferences** page enables you to configure everRun 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 66 .
Product License	Allows you to view and manage the everRun product license. See "Managing the everRun Product License" on page 67 .
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 71 .
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 environment. See "Quorum Servers" on page 14 and "Configuring Quorum Servers" on page 73 .
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 75 .
System Resources	Allows you to specify the number of virtual CPUs (VCPUs) and the amount of memory reserved for the everRun software. See "Configuring System Resources" on page 76 .
Administrative Tools	

Preference	Description
Active Directory	Allows you to enable (and disable) Active Directory. When enabled, Active Directory allows you to authorize existing users or groups from the Active Directory domain to log on to the everRun Availability Console and manage the everRun system. See "Configuring Active Directory" on page 77 .
Migration Policy	Allows you to disable VMs' automatic load balancing that occurs, by default, when the node they are running on returns to service after recovering from a node failure or exiting maintenance mode. See "Configuring the Migration Policy" on page 79 .
Secure Connection	Allows you to enable only HTTPS connections to the system. See "Configuring Secure Connections" on page 79 .
Host Inactivity Logout	Allows you to disable the host inactivity logout or to change the timeout. See "Configuring the Host Inactivity Logout" on page 80 .
Snapshot Configuration	Allows you to disable the creation of snapshots. See "Disabling and Enabling Snapshots" on page 81 .
IPtables Security	Allows you to manage IP packet filtering using the administrative tool IPtables.
Diagnostics	
Diagnostics	Allows you to generate a diagnostic file for your authorized Stratus service representative. See "Managing Diagnostic Files" on page 85 .
Notification	
e-Alerts	Allows you to enable email alerts (e-Alerts) for system administrators. See "Configuring e-Alerts" on page 87 .
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 89.
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 system can send notifications to your authorized Stratus service representative about problems with the system. See "Configuring Remote Support Settings" on page 91.
Proxy Configuration	Allows you to configure proxy settings for the everRun 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 93.

Related Topics

["The everRun Availability Console" on page 58](#)

["Using the everRun Availability Console" on page 57](#)

Specifying Owner Information

Specify the name and contact information for an administrator or owner of the everRun 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**.

Related Topics

["The Preferences Page" on page 63](#)

["The everRun Availability Console" on page 58](#)

Managing the everRun Product License

Manage the product license for the everRun system by:

- Viewing current license information such as status, type, expiration date, and other information.
- 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 system.
- Activating, renewing, or checking the status of an existing license.

When you purchase an everRun 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 system) that you can access when you need to upload (and activate) the license to the everRun 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 Customer Support or your authorized Stratus service representative. See the **everRun Support** page at <http://www.stratus.com/services-support/customer-support>.

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

Follow the procedure below that is appropriate for your needs.

To upload a new license .key file to an everRun 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 system. The everRun system must have Internet connectivity.

1. Click **Preferences** in the left-hand navigation panel, to open the **Preference** page.
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 **Choose File** (in Google Chrome) or **Browse** (in Firefox or Internet Explorer), 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 system. The everRun system contacts the Stratus server to activate the license.

To license an everRun system with no Internet connectivity but with a connection to a computer that has Internet connectivity

If your everRun 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 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 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 **Choose File** and navigate to the license .key file that you saved in the previous step. Then, click **Upload** to upload it to the everRun system.

To license an everRun system with no Internet connectivity

If your everRun 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 system.

For this procedure:

- You will need a USB flash drive and two computers (A and B) in addition to the everRun system.
- Computer A has Internet access and has no connection to the everRun system.
- Computer B has access to the everRun Availability Console on the everRun system, but Computer B and the everRun system 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 is automatically 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 **Choose File**, 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 <i>nn</i> days <i>nn</i> hours
TYPE OF LICENSE	Enterprise Edition (volume)
EXPIRATION	<i>month dd, 20yy, time</i>
LAST CHECK	<i>month dd, 20yy, time</i>
ASSET ID	<i>month dd, 20yy, time</i>
FT Enabled	<i>true_or_false</i>
Split Site Allowed	<i>true_or_false</i>

If a license activation fails, the License Activation Server (or ALAS) returns a numeric error code. The following drop-down menu lists the error codes.

To view the 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 58](#)

["The Preferences Page" on page 63](#)

["Using the everRun Availability Console" on page 57](#)

Configuring IP Settings

Configure Internet Protocol (IP) settings for the everRun 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 software, you configure three IP addresses: one for the everRun system and one for each node. You can change the IP addresses and other IP settings after

installation using the appropriate procedure below. You must specify a static IPv4 address for the everRun system.

Warnings:



1. Do not change the IP configuration settings, especially on systems with running VMs, without the advice and knowledge of your network administrator. Doing so could make the system and all its VMs inaccessible.
2. If you change the **Static System IP** address, any MAC addresses automatically assigned to the VMs will change when the VMs reboot, because the everRun software generates MAC addresses for the VMs based on the system IP address. To prevent changes to the MAC address for a VM (for example, to support software applications that are licensed on a MAC-address basis), set a persistent MAC address as described in "[Assigning a Specific MAC Address to a Virtual Machine](#)" on page 263.
3. You must use the everRun Availability Console to change IP addresses. Do not use Linux tools.

Notes:



1. The procedure you use to configure IP settings depends on whether the everRun system stays on the same subnet or moves to a new subnet. For instructions on how to move the system to a new subnet, access the Knowledge Base to search for the article *Moving an everRun System to a Different Subnet* (KB-4264). See "[Accessing Knowledge Base Articles](#)" on page 481. Follow the appropriate procedure for your needs.
2. Changing IP settings for a new subnet typically includes changing the node's physical network connections (for example, disconnecting and then re-attaching network cables if moving the PMs). Before you disconnect cables from nodes, you must shut down the nodes.
3. In a single-node system, the **IP Configuration** page displays settings for only one node.

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

The everRun 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. In the **Static System IP** box, type the static system IP address that you obtained from your network administrator.
4. Click the **Static** button and type valid, unique values for **Primary DNS** and **Secondary DNS**.
5. Verify that the displayed **NetMask** value is correct.
6. For **Node0** and **Node1**, enter appropriate values for **IP Address** and **Gateway IP**.
7. Click **Save** to save the values (or click **Reset** to restore previous values).

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

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

Related Topics

["Software Installation" on page 34](#)

["Obtaining System IP Information" on page 53](#)

["Logging on to the everRun Availability Console for the First Time" on page 54](#)

["The Preferences Page" on page 63](#)

["The everRun Availability Console" on page 58](#)

["Using the everRun Availability Console" on page 57](#)

Configuring Quorum Servers

When you log on to the everRun 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 32](#).

Notes:



1. For a VM to recognize quorum server configuration changes, you must reboot the VM by shutting it down and then restarting it. See ["Shutting Down a Virtual Machine" on page 225](#) and ["Starting a Virtual Machine" on page 225](#).
2. Windows Updates on a quorum server can interrupt the server's operation, which affects fault-recovery behavior. On quorum servers, you should schedule Windows Updates during a maintenance period or disable Windows Updates.

To configure quorum servers

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

Click **Save** to save the values.

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

To remove a quorum server



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

1. Navigate to the **Preferences** page of the 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 58](#)

["The Preferences Page" on page 63](#)

["Using the everRun Availability Console" on page 57](#)

Configuring Date and Time

When you log on to the everRun 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. Click **Preferences** in the left-hand navigation panel, to open the **Preferences** page.
2. On the **Preferences** page, click **Date & Time**.

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



Note: If you configure time manually, the 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 Console masthead 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 58](#)

["The Preferences Page" on page 63](#)

["Using the everRun Availability Console" on page 57](#)

Configuring System Resources

Configure system resources to specify how the everRun 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 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 software. Values are **2** (the default) and **4**.
 - **System Memory**, which sets the amount of memory reserved for the everRun 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 58](#)

["The Preferences Page" on page 63](#)

["Using the everRun Availability Console" on page 57](#)

Configuring Active Directory

Configure Active Directory for the everRun system to authorize existing users or groups from an Active Directory domain to log on to the everRun Availability Console with their Active Directory credentials.

*After you add the everRun system to an Active Directory domain, you can assign administrative privileges to domain users using the **Grant Access** wizard, which you start from the **Users & Groups** page (see ["The Users & Groups Page" on page 109](#)).*

To add the everRun system to an Active Directory domain

1. Click **Preferences** in the left-hand navigation panel, to open the **Preferences** page.
2. Click **Active Directory**.
3. Click **Enable Active Directory**.
4. Next to **Active Directory Domain**, type the name of the domain to use.
5. Click one of the following to select whether the system will prevent or allow all AD users from being automatically assigned the "everyone" role:
 - **Prevent all AD users from being automatically assigned the "Everyone" role** (the default).
 - **Allow all AD users to authenticate and be authorized for "Everyone" role access.**
6. Click **Add System to Active Directory**.

7. Type the **Username** and **Password** of an Active Directory Administrator in order to add this everRun system to the domain.
8. Click **Add**.
9. Assign administrative privileges to domain users on the **Users & Groups** page, as described in ["Managing Domain User Accounts" on page 112](#).

To remove an everRun system from an Active Directory domain

1. In the everRun Availability Console, click **Preferences** in the left panel, to open the **Preferences** page.
2. Click **Active Directory**.
3. Click **Remove System from Active Directory**.
4. Type a **Username** and **Password** that provides you with administrative privileges within the domain.
5. Click **Remove**.

To disable domain authentication

1. In the everRun Availability Console, click **Preferences** in the left panel, to open the **Preferences** page.
2. Click **Active Directory**.
3. Click **Disable Active Directory**.



Note: Disabling Active Directory prevents the use of domain authentication for authorizing administrators of the everRun system; however, it does not remove the system from the domain. To restore the use of domain authentication, click **Enable Active Directory**. You do not need to retype the name of the controller or restore domain users on the **Users & Groups** page.

Related Topics

["The Users & Groups Page" on page 109](#)

["Managing Domain User Accounts" on page 112](#)

["Managing Local User Accounts" on page 110](#)

["The Preferences Page" on page 63](#)

["The everRun Availability Console" on page 58](#)

Configuring the Migration Policy

By default, VMs automatically load balance when the node they are running on returns to service after recovering from a node failure or exiting maintenance mode. You can disable this automatic load balancing by setting the migration policy.

To set the migration policy

1. Click **Preferences** in the left-hand navigation panel, to open the **Preference** page.
2. On the **Preferences** page, click **Migration Policy**.
3. Select **Disable automatic load balancing after bringing a node back into service** to prevent VMs from automatically load balancing.
4. Click **Save**.

After the migration policy is set and a node returns to service, a load-balancing scale () appears in the masthead with the message **VMs are not load balanced** and a link to ["Load Balancing" on page 129](#). Click the link to rebalance the load.

Related Topics

["Managing Virtual Machines" on page 135](#)

["The Preferences Page" on page 63](#)

["The everRun Availability Console" on page 58](#)

Configuring Secure Connections

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

Note:



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

To enable only HTTPS connections

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

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

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

To enable HTTP and HTTPS connections

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

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

Related Topics

["The everRun Availability Console" on page 58](#)

["The Preferences Page" on page 63](#)

["Using the everRun Availability Console" on page 57](#)

Configuring the Host Inactivity Logout

For security, an everRun system limits the inactivity time of a login session on a host operating system. The default timeout is 10 minutes. After 10 minutes (or other specified time) of inactivity, the everRun system automatically logs out the session. A host inactivity logout prevents a login session from remaining open indefinitely without use.

To enable the Host Inactivity Logout and set the timeout

1. Click **Preferences** in the left-hand navigation panel, to open the **Preferences** page.
2. On the **Preferences** page, click **Host Inactivity Logout**.
3. Activate the check box next to **Enable Host Inactivity Logout**.
4. The default **Host Inactivity Logout** timeout is 10 minutes. To specify a different timeout, type a number of minutes next to **Timeout Minutes**.

Enter minutes in whole numbers. You cannot enter 0.
5. Click **Save** to save the values (or click **Reset** to restore previous values).

Related Topics

["The Preferences Page" on page 63](#)

["The everRun Availability Console" on page 58](#)

Disabling and Enabling Snapshots

Snapshots provide an image of a VM at a particular moment in time. By default, the everRun system's ability to take snapshots is enabled. At times, you may want to disable the system's ability to create snapshots, for security reasons. Or, if disabled, you may want to re-enable the system's ability to create snapshots.

To disable the ability to take snapshots

1. Click **Preferences** in the left-hand navigation panel, to open the **Preference** page.
2. On the **Preferences** page, click **Snapshot Configuration**.
3. Activate the check box next to **Disable Snapshot**.
4. Click **Save**.

If snapshots are disabled and you want to take a snapshot, you need to enable snapshots.

To enable the ability to take snapshots

1. Click **Preferences** in the left-hand navigation panel, to open the **Preference** page.
2. On the **Preferences** page, click **Snapshot Configuration**.
3. Deactivate the check box next to **Disable Snapshot**.
4. Click **Save**.

Related Topics

["The Snapshots Page" on page 104](#)

["Managing Snapshots" on page 249](#)

["Using the everRun Availability Console" on page 57](#)

Managing IPtables

The administration tool for managing IP packet filtering for the Linux operating system is known as *iptables*. With everRun systems, the task of working with iptables has been simplified and streamlined. Using the **IPtables Security** page, you can set up, maintain, and inspect the various filter table chains and their underlying rules. You have access to the three main chains (**INPUT**, **OUTPUT**, and **FORWARD**) for applying the packet-filtering rules you need. With everRun systems, the rules are applied to each physical machine (PM), to both IPv4 and IPv6 packets, and the rules remain persistent after rebooting.

When you create a rule, you specify a chain (**INPUT**, **OUTPUT**, or **FORWARD**) and a **Rule ID**. When processing inbound packets, the kernel applies the rules associated with the **INPUT** chain, and when processing outbound packets, the kernel applies the rules associated with the **OUTPUT** chain. The kernel applies the rules associated with the **FORWARD** chain when processing received inbound packets that must be routed to another host. Rules are applied in order of the **Rule ID**. (A **Rule ID** is similar to a row ID, where, for example, **Rule ID** 1 equals row 1.) Instead of creating rules, however, you can load default settings for the rules.

The **IPtables Security** page displays a separate table for each of the three chains and their associated rules. The rules, if they exist for a particular chain, are sorted by **Rule ID**. Columns display the network name, type of network, protocol, and other information. If necessary, use the scroll-bar on the right side of the page to view all of the rules and the scroll-bar at the bottom to view all of the columns. For more information on iptables functionality, see the Linux manual (man) pages for iptables.

Notes:



1. For information on the ports that everRun software uses, see ["System Requirements Overview" on page 24](#).
2. For additional information on everRun TCP and UDP ports, access the Knowledge Base to search for the article *TCP and UDP ports used by everRun 7* (KB-2123). See ["Accessing Knowledge Base Articles" on page 481](#).

To manage IPtables, perform the appropriate task below.

To create a new rule

1. Click **Preferences** in the left-hand navigation panel, to open the **Preference** page.
2. On the **Preferences** page, click **IPtables Security**.
3. Ensure that the **Enable Port Management** box has a check mark.
4. Click the **Insert New Rule** button to open the **Insert New Rule** pop-up window.
5. In the **Insert New Rule** pop-up window, set values for the following:
 - **Chain**—Select **INPUT**, **OUTPUT**, or **FORWARD** from the drop-down list.
 - **Rule ID**—Enter a number that establishes the order for processing the rule. Enter a value, starting with 1 and up to a maximum value that is the total number of rules within the chain. Each **Rule ID** value must be unique.
If you enter a number that is already assigned to a rule, the existing rule is incremented by 1 (as are subsequent rules, if any) and the number you enter is assigned to the new rule. So, if, for example, **Rule ID 1** already exists and you enter 1 for the new rule, the existing **Rule ID 1** becomes **Rule ID 2**, the existing **Rule ID 2** (if it exists) becomes **Rule ID 3**, and so on.
 - **Shared Network**—Select a network from the drop-down list of all available shared networks.
 - **Protocol**—Select **udp**, **tcp**, or **all**.
Selecting **all** causes the **Grouping** and **Port Number** fields to become inactive (gray) because setting a range of port numbers is unnecessary.
 - **Target**—Select **drop**, **accept**, or **reject** for the action you want to apply to the packet that matches the rule's specifications.
 - **Grouping**—Select one of the following values:
 - **range**: If you select this value, both **Port Number** fields (see below) become active and you must enter values in the fields.
 - **single**: If you select this value, the **Port Number (starting)** field becomes active, and you must enter a value in it.

- **all**: Selecting this value causes both **Port Number** fields to become inactive (gray). (The value **all** for **Grouping** is selected automatically when you select **all** for **Protocol**.)
- **Port Number (starting)**—Enter a number 0 to 65535 for the first port of the range.
- **Port Number (ending)**—Enter a number 1 to 65535 that is greater than **Port Number (starting)** for the last port of the range.

Click **Insert** to insert the new rule.

6. Click **Save** at the bottom of the page, or click **Reset** to cancel any unsaved changes, which restores rules to those of the last saved session.

After the new rule is saved, the **IPtables Security** page displays it in the appropriate chain.

To remove a rule

1. Click **Preferences** in the left-hand navigation panel, to open the **Preference** page.
2. On the **Preferences** page, click **IPtables Security**.
3. Ensure that the **Enable Port Management** box has a check mark.
4. Select the rule that you want to remove.
5. Click **Remove** (in the right-most column), for the rule you selected.
6. Click **Save** at the bottom of the page, or click **Reset** to cancel any unsaved changes, which restores rules to those of the last saved session.

After the rule is removed, it disappears from the **IPtables Security** page .

To load default settings



Caution: Loading default settings will override current settings. .

1. Click **Preferences** in the left-hand navigation panel, to open the **Preference** page.
2. On the **Preferences** page, click **IPtables Security**.
3. Ensure that the **Enable Port Management** box has a check mark.
4. Click **Load Default Settings** at the bottom of the page.

A warning appears: *Current settings will be overridden by the initial settings!* Click **OK** if you want to load the default settings, or click **Cancel** to cancel the loading of default settings.

Related Topics

["The Preferences Page" on page 63](#)

["The everRun Availability Console" on page 58](#)

Managing Diagnostic Files

Diagnostic files provide a snapshot of an everRun 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 85](#)
- ["Deleting a Diagnostic File" on page 87](#)
- ["Uploading a Diagnostic File to Customer Support" on page 86](#)

Related Topics

["The everRun Availability Console" on page 58](#)

["Using the everRun Availability Console" on page 57](#)

["The Preferences Page" on page 63](#)

Creating a Diagnostic File

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



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

To create diagnostic files

1. Click **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.
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 86.

Related Topics

["The everRun Availability Console" on page 58](#)

["The Preferences Page" on page 63](#)

["Using the everRun Availability Console" on page 57](#)

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 system. (To create a diagnostic file, see "[Creating a Diagnostic File](#)" on page 85.)

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 system has Internet connectivity, upload the diagnostic file directly to the Stratus everRun Customer Support web site by clicking **Upload**. An **Alert** appears with the message that the diagnostic was uploaded successfully, if the upload succeeds.
 - If the everRun 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 **Choose File**, select the zip file on the computer, and click **Submit**.

If you need help with this procedure, call everRun Customer Support at the phone number listed on the **everRun Support** page at <http://www.stratus.com/services-support/customer-support>.

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 system, as described in "[Deleting a Diagnostic File](#)" on page 87.

Related Topics

["The everRun Availability Console" on page 58](#)

["The Preferences Page" on page 63](#)

["Using the everRun Availability Console" on page 57](#)

Deleting a Diagnostic File

Delete a diagnostic file from the everRun 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 58](#)

["The Preferences Page" on page 63](#)

["Using the everRun Availability Console" on page 57](#)

Configuring e-Alerts

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

To enable e-Alerts

1. Click **Preferences** in the left-hand navigation panel, to open the **Preference** page.
2. Under **Notification**, click **e-Alerts**.
3. Click the **Enable e-Alerts** box. Boxes for specifying or selecting the following settings appear:

- **SMTP Server** (required)—Enter the name of the Simple Mail Transfer Protocol (SMTP) server that your company uses to send email.
- **Port Number** (optional)—Enter the port number to use when sending e-Alerts. If no port number is specified, the default SMTP port 25 will be used. (For additional information on all ports, including the SMTP port, access the Knowledge Base to search for the article *TCP and UDP ports used by everRun 7* (KB-2123). See "[Accessing Knowledge Base Articles](#)" on [page 481](#).)
- **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.

- **Encrypted Connection**—Select a value from the pull-down menu for the encryption protocol that the SMTP server requires:
 - **None** for no encryption. By default, port number 25 is used.
 - **TLS** for the Transport Layer Security (TLS) protocol. For TLS, Stratus recommends that you specify 587 for **Port Number**, though 25 is used by default.
 - **SSL** for the Secure Sockets Layer (SSL) protocol. For SSL, Stratus recommends that you specify 465 for **Port Number**, though 25 is used by default.
- **Enable Authentication**—Click this box if the SMTP server requires authentication to send email 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 and sends a sample email with the subject "Test Alert" to all email recipients; SNMP sends traps to recipients of SNMP traps, if configured (see "[Configuring SNMP Settings](#)" on page 89); and Support Configuration sends a notification to your authorized Stratus service representative, if configured (see "[Configuring Remote Support Settings](#)" on page 91). Watch the Alerts History log (see "[The Alerts Page](#)" on page 94) for delivery status.

You can also test e-Alerts by putting the secondary physical machine into maintenance mode (see "[Maintenance Mode](#)" on page 125), 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 58](#)

["The Preferences Page" on page 63](#)

["Using the everRun Availability Console" on page 57](#)

Configuring SNMP Settings

Configure Simple Network Management Protocol (SNMP) settings for your everRun 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. You can download a copy of the MIB file from the **Drivers and Tools** section of the **Downloads** page at <http://www.stratus.com/go/support/everrun>.
- **SNMP trap**—a message initiated by one of the nodes in 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
```

Follow the appropriate procedure to enable SNMP requests or traps.

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



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

1. Click **Preferences** in the left-hand navigation panel, to open the **Preference** page.
2. Under **Notification**, click **SNMP Configuration**.
3. Activate the check box next to **Enable SNMP Traps**.
4. Type the name of the **SNMP Community**, or keep the default (**public**).
5. Next to **List of Recipients for SNMP traps (Version 2c)**, type the IP address or host name for each SNMP version 2 recipient, one per line.
6. Next to **List of Recipients for SNMP traps (Version 1)**, type the IP address or host name for each SNMP version 1 recipient, one per line.
7. Click **Save**. (Or click **Reset** to restore the previously saved values.)

8. Configure your organization's firewall to allow SNMP operations, which enables SNMP management systems to receive alerts from and send traps to the everRun system. To do so, configure your organization's firewall to open the SNMP port:

Message Type: SNMP

Protocol: SNMP

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

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

The everRun software generates a test alert and SNMP sends traps to recipients of SNMP traps; e-Alerts send a sample email with the subject "Test Alert" to all email recipients of e-Alerts, if configured (see ["Configuring e-Alerts" on page 87](#)); and Support Configuration sends a notification to your authorized Stratus service representative, if configured (see ["Configuring Remote Support Settings" on page 91](#)). Watch the Alerts History log (see ["The Alerts Page" on page 94](#)) for delivery status.

Related Topics

["SNMP" on page 484](#)

["The everRun Availability Console" on page 58](#)

["The Preferences Page" on page 63](#)

["Using the everRun Availability Console" on page 57](#)

Configuring Remote Support Settings

When you log on to the everRun 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



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.

1. Click **Preferences** in the left-hand navigation panel, to open the **Preferences** page.
2. Under **Remote Support**, click **Support Configuration**.

3. Modify the settings, as appropriate for your system:

- **Enable Remote Support Access** allows your authorized Stratus service representative to remotely connect to the everRun system for troubleshooting purposes. Note that you can enable and then disable this 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.

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

5. Configure your organization's firewall to allow support messages.

To configure your firewall to allow support messages

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

Message Type: Call-Home and Licensing

Protocol: TCP

Port: 443

Stratus support server address: *.stratus.com

Message Type: Support Diagnostics

Protocol: TCP

Port: 443

Stratus support server address: *.stratus.com

Message Type: Dial-In

Protocol: TCP

Port: 443, Default proxy port: 3128 (You can change the default proxy port number.)

Stratus support server address: *.ecacsupport.com

Message Type: e-Alert

Protocol: SMTP

Port: 25

(For additional information on TCP and UDP ports, access the Knowledge Base to search for the article *TCP and UDP ports used by everRun 7* (KB-2123). See ["Accessing Knowledge Base Articles" on page 481.](#))

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

Message Type: SNMP

Protocol: SNMP

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

6. Generate a test alert.

To generate a test alert

Click **Generate Test Alert**. The everRun software generates a test alert and Support Configuration sends a notification to your authorized Stratus service representative; e-Alerts send a sample email with the subject "Test Alert" to all email recipients of e-Alerts, if configured (see ["Configuring e-Alerts" on page 87](#)); and SNMP sends traps to recipients of SNMP traps, if configured (see ["Configuring SNMP Settings" on page 89](#)). Watch the Alerts History log (see ["The Alerts Page" on page 94](#)) for delivery status. A subsequent alert will be generated if the support notification fails.

Related Topics

["The everRun Availability Console" on page 58](#)

["The Preferences Page" on page 63](#)

["Using the everRun Availability Console" on page 57](#)

Configuring Internet Proxy Settings

Configure proxy settings for the everRun 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 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 58](#)

["The Preferences Page" on page 63](#)

["Using the everRun Availability Console" on page 57](#)

The Alerts Page

The **Alerts** page displays messages about events on the everRun 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 95.](#))

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. (You can also display alert information using `snmptable`; see ["Obtaining System Information with snmptable" on page 484.](#))

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 91](#)
- ["Configuring e-Alerts" on page 87](#)
- ["Configuring SNMP Settings" on page 89](#)

Related Topics

["The everRun Availability Console" on page 58](#)

["Using the everRun Availability Console" on page 57](#)

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 94.](#))

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.

You can also display information about audit logs using `snmptable` (see ["Obtaining System Information with snmptable" on page 484.](#))

Related Topics

["The everRun Availability Console" on page 58](#)

["Using the everRun Availability Console" on page 57](#)

The Physical Machines Page

The **Physical Machines** page enables you to manage the physical machines (PMs) in the everRun 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 98.](#) To display information about a node, you can use the `snmptable` command; see ["Obtaining System Information with snmptable" on page 484.](#)

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

- Action buttons: Various action buttons appear, with inactive buttons grayed out, depending upon the state of the selected node. 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 125](#)). To learn about additional PM actions available in maintenance mode, see ["Physical Machine Actions" on page 96](#) 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 58](#)

["Using the everRun Availability Console" on page 57](#)

Physical Machine Actions

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



Caution: Use the **Physical Machines** page of the 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
 <p>Work On</p>	<p>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 125.</p>
<p>The following actions are available after clicking the Work On button, when the PM has entered maintenance mode.</p>	
 <p>Finalize</p>	<p>Removes a PM from the state running (in Maintenance). See "Maintenance Mode" on page 125.</p>
 <p>Shutdown</p>	<p>Shuts down a PM. The PM transitions to off (in Maintenance). See "Shutting Down a Physical Machine" on page 127.</p>
 <p>Reboot</p>	<p>Reboots the PM. The PM transitions to preparing for reboot (in Maintenance). See "Rebooting a Physical Machine" on page 127.</p>
 <p>Remove</p>	<p>Causes the everRun 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 274.</p>
<p>The following actions may be available when a PM has failed or when everRun software has removed a PM from service and powered it off, due to an excessive failure rate.</p>	
 <p>Reset Device</p>	<p>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 132.</p>
 <p>Recover</p>	<p>Recovers a failed PM. In some cases, the everRun Availability Console displays the state of a failed PM as Unreachable (Syncing/Evacuating...). See "Recovering a Failed Physical Machine" on page 130.</p>

Related Topics

["The everRun Availability Console" on page 58](#)

["Using the everRun Availability Console" on page 57](#)

["The Physical Machines Page" on page 95](#)

Physical Machine States and Activities

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

State	Activity	Available Commands	Description
	 Running	Work On	PM is running normally.
	 Evacuating	Finalize	Virtual machines are migrating from this PM to its partner.
	 Running	Work On	PM is predicted to fail.
	 Running	Work On	PM failed.
	 Powered Off	Work On 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 132 .
	 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 125 .

Related Topics

["The everRun Availability Console" on page 58](#)

["Using the everRun Availability Console" on page 57](#)

["The Physical Machines Page" on page 95](#)

The Virtual Machines Page

Use the **Virtual Machines** page to manage the virtual machines (VMs) running on your everRun 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 102](#). To learn more about the controls on this page, see ["Virtual Machine Actions" on page 100](#) or the help topic for a specific task.

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

- [Creating, copying, exporting, importing, or restoring VMs, as described in "Creating and Migrating Virtual Machines" on page 141](#)
- ["Opening a Virtual Machine Console Session" on page 227](#)
- ["Reprovisioning Virtual Machine Resources" on page 231](#)
- [Creating VM snapshots that can be restored or exported, as described in "Creating a Snapshot" on page 250](#)
- [Controlling the power state of a VM, as described in:](#)
 - ["Starting a Virtual Machine" on page 225](#)
 - ["Shutting Down a Virtual Machine" on page 225](#)
 - ["Powering Off a Virtual Machine" on page 226](#)
- ["Removing a Virtual Machine" on page 229](#) or ["Renaming a Virtual Machine" on page 228](#)
- [Performing advanced tasks or troubleshooting, as summarized in "Advanced Topics \(Virtual Machines\)" on page 263](#)

- Viewing information about a VM, including its name, operating system, description, and resources in the tabs of the bottom pane

Related Topics

["Managing Virtual Machines" on page 135](#)

["Using the everRun Availability Console" on page 57](#)

Virtual Machine Actions

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

Action	Description
 Create	Launches the Create VM Wizard. See "Creating a New Virtual Machine" on page 143 .
 Copy	Copies an existing VM on your system to create a new VM or to create a duplicate VM for troubleshooting. See "Copying a Virtual Machine" on page 149 .
 Import/Restore	<p>Imports a VM from a set of OVF and VHD files. See "Creating and Migrating Virtual Machines" on page 141.</p> <p>The import wizard allows you to <i>import</i> a VM to create a new instance of the VM or <i>restore</i> a VM to create an identical VM with the same hardware IDs provided in the OVF and VHD files.</p> <p>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.</p>
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 227 .

Action	Description
 Snapshot	Creates a VM snapshot that you can export to OVF and VHD files. See "Managing Snapshots" on page 249 .
 Shutdown	Shuts down the selected VM. See "Shutting Down a Virtual Machine" on page 225 .
 Power Off	Immediately stops processing in the selected VM and destroys its memory state. Use this only as a last resort, when the VM cannot be successfully shutdown. See "Powering Off a Virtual Machine" on page 226 .
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 231
 Export	Saves the image of a VM to a set of OVF and VHD files. You can import these files on another system or import them back to the same everRun system to restore or duplicate the original VM. See "Exporting a Virtual Machine from an ever-Run System" on page 208 .
 Restore	Recovers an existing VM on your everRun system by overwriting the VM from a previous backup copy of OVF and VHD files. See "Replacing a Virtual Machine from an OVF File" on page 203 .
 Snapshot	Creates a VM snapshot that you can use to create a new VM or export to OVF and VHD files. See "Managing Snapshots" on page 249 .
 Start	Boots the selected VM. See "Starting a Virtual Machine" on page 225 .

Action	Description
 Boot From CD	Boots a VM from the selected virtual CD. See "Booting from a Virtual CD" on page 247 .
The following action is available if the everRun software has removed the VM from service and powered it off because an excessive failure rate.	
 Reset Device	Resets the mean time between failures (MTBF) counter for a VM so it can be brought back into service. See "Resetting MTBF for a Failed Virtual Machine" on page 267 . 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.

Related Topics

["Managing the Operation of a Virtual Machine" on page 224](#)

["The Virtual Machines Page" on page 99](#)

["Using the everRun Availability Console" on page 57](#)

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 software is installing the boot volume for a new VM.
	 stopped	Start Copy	VM has been shut down or powered off.

State	Activity	Enabled Actions	Description
		Config Export Snapshot Boot From CD Remove	
	 booting	Console Power Off	VM is starting.
	 running	Console Snapshot 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 messages are sent.
	 crashed		VM crashed too many times and exceeded its MTBF threshold. The VM is left in a crashed state until Reset Device is clicked. See " Resetting

State	Activity	Enabled Actions	Description
			MTBF for a Failed Virtual Machine" on page 267.

Related Topics

["Managing the Operation of a Virtual Machine" on page 224](#)

["The Virtual Machines Page" on page 99](#)

["Using the everRun Availability Console" on page 57](#)

The Snapshots Page

Use the **Snapshots** page to manage virtual machine (VM) snapshots, which represent an image of a VM at a particular point in time. You can use a snapshot to restore a VM on the everRun system or you can export a snapshot for use in a new VM. To open this page, click **Snapshots** in the left-hand navigation panel of the everRun Availability Console.

To create a snapshot (on the **Virtual Machines** page), see ["Creating a Snapshot" on page 250](#).

The everRun system's ability to take snapshots is enabled, by default. To disable or to re-enable the system's ability to take snapshots, see ["Disabling and Enabling Snapshots" on page 81](#).

To manage an existing snapshot, click the name of a snapshot in the top pane of the **Snapshots** page. The bottom pane displays a description of the snapshot.

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

- ["Exporting a Snapshot" on page 256](#)
- ["Creating a Virtual Machine from a Snapshot" on page 254](#)
- ["Removing a Snapshot" on page 262](#)
- Adding a description for each volume, in the **Description** text box

Related Topics

["Managing Snapshots" on page 249](#)

["Using the everRun Availability Console" on page 57](#)

The Volumes Page

The **Volumes** page displays information about volumes that are attached to virtual machines (VMs) in the everRun 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:

- **State**
- **Name**
- **Disk Synchronization**
- **Sector Size**
- **Size**
- **Bootable**
- **Storage Group**
- **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 also display information about volumes using the `snmptable` command; see ["Obtaining System Information with snmptable" on page 484](#).) 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 240](#)).
- View information about the volume container, including the volumes and snapshots it contains, on the **Container** tab.
- Expand a volume container on the **Container** tab (see ["Expanding a Volume Container on the everRun System" on page 241](#)).

- Remove a volume by clicking **Remove**. Note, though, that the **Remove** button is grayed out when a VM is using a volume.

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

- ["Attaching a Volume to a Virtual Machine" on page 235](#)
- ["Creating a Volume in a Virtual Machine" on page 234](#)
- ["Detaching a Volume from a Virtual Machine" on page 237](#)
- ["Removing a Volume from a Virtual Machine" on page 238](#)

Related Topics

["The everRun Availability Console" on page 58](#)

["Using the everRun Availability Console" on page 57](#)

The Storage Groups Page

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

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

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

Columns in the **Summary** tab display information about the logical ID, size, sector size, and state of each disk within the group as well as the PM on which the disk runs. To display or hide columns, move the cursor to the right of a column heading, click the down-arrow that appears, and then click **Columns**, selecting or de-selecting the columns that you want to show or hide.



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 Storage Architecture" on page 15](#).

Related Topics

["The everRun Availability Console" on page 58](#)

["Using the everRun Availability Console" on page 57](#)

The Networks Page

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

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

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

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

- ["Connecting Additional Networks" on page 55](#).
- ["Fixing a Network Connection" on page 108](#).
- 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 27](#)
- ["SplitSite Network Requirements" on page 30](#)



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 58](#)

["Using the everRun Availability Console" on page 57](#)

Fixing a Network Connection

The everRun system software monitors and analyzes network connections. If it determines that an existing network connection is not optimal (for example if a 1Gb port is connected to a 10Gb port), and it cannot automatically reconfigure the network, it issues an alert stating that cabled network ports could not be paired automatically. In this case, perform the following procedure to reconfigure network connections so that they are optimal.

To reconfigure non-optimal network connections

1. Place the secondary PM into maintenance mode. See ["Maintenance Mode" on page 125](#) for details.
2. Open the **Networks** page in the everRun Availability Console.
3. Click the **Fix Network** button if it is active (the button is inactive if networks have no problems or no fixable problems). As the everRun system software reconfigures the networks, the connection topology displayed in the diagram on the **Networks** page will change to show the new optimal configuration.
4. Remove the secondary PM from maintenance mode. See ["Maintenance Mode" on page 125](#) for details.

Related Topics

["The Networks Page" on page 107](#)

["The everRun Availability Console" on page 58](#)

["Using the everRun Availability Console" on page 57](#)

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 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 244](#)
- ["Removing a Virtual CD" on page 248](#)
- ["Renaming a Virtual CD" on page 248](#)
- Adding a description for each volume, in the **Description** text box

To complete other VCD management tasks, see ["Managing Virtual CDs" on page 244](#).

Related Topics

["Inserting a Virtual CD" on page 246](#)

["Ejecting a Virtual CD" on page 246](#)

["Using the everRun Availability Console" on page 57](#)

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. 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 Software" on page 115](#).

Related Topics

["The everRun Availability Console" on page 58](#)

["Using the everRun Availability Console" on page 57](#)

The Users & Groups Page

Use the **Users & Groups** page to add, modify, or remove user accounts on, or to grant access for Active Directory users to manage 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 110](#).

To manage domain user accounts

For information about enabling the Active Directory service on your everRun system, see ["Configuring Active Directory" on page 77](#). To grant or remove access for domain users to manage the everRun system, see ["Managing Domain User Accounts" on page 112](#).



Note: If you are logged on as administrator to a system that has Active Directory users or groups configured, the **Grant Access** button will appear in the upper-right corner of the **Users & Groups** page. Clicking the **Grant Access** button launches the Grant Access wizard. The ["Managing Domain User Accounts" on page 112](#) topic discusses using the Grant Access wizard.

To sort and locate user accounts

If you have a large number of accounts, you can click a column heading to sort the accounts by parameter. You can sort accounts by **Type**, **Username**, **Real Name**, **Email** address, or **Role**.

Related Topics

["Configuring Active Directory" on page 77](#)

["Managing Domain User Accounts" on page 112](#)

["Managing Local User Accounts" on page 110](#)

Managing Local User Accounts

You add, edit, or remove users, specify passwords, and assign user roles to local-user accounts on the **User & Groups** page in the everRun Availability Console. (To grant or deny access for established user accounts in an Active Directory domain, see ["Managing Domain User Accounts" on page 112](#).)

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.

User roles are:

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

To add a user account

1. In the left-hand navigation 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** (and **Confirm Password**), **Email Address**, and **Real Name** fields. User names and passwords may be from 1 to 64 characters long, and must include no white space.
5. Click **Save**.

To edit a user account

1. In the left-hand navigation panel, select **Users & Groups**.
2. Under **USERS & GROUPS**, select the account you want to edit.
3. In the top pane, click **Edit**.
4. To change a user's role, in the **Role** drop-down window, select **Administrator**, **Platform Manager**, or **Read-only**.
5. Click **Save**.

To remove a user account

1. Under **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.

Related Topics

["Configuring Active Directory" on page 77](#)

["Managing Domain User Accounts" on page 112](#)

["The Users & Groups Page" on page 109](#)

Managing Domain User Accounts

You can grant Active Directory (AD) domain user accounts access to the everRun Availability Console. Domain user accounts are managed on a central AD domain server, as opposed to the local everRun system.

After granting access to domain accounts, you can use the Grant Access wizard (on the Users & Groups page) to view, manage, and sort the AD accounts that have access to the system.



Prerequisites: You must add the everRun system to an Active Directory domain before you can manage domain accounts. (See ["Configuring Active Directory" on page 77.](#)) If Active Directory is not configured, or if the user who is logged onto the interface does not have administrator privileges, the Grant Access button is grayed out on the Users & Groups page.

To grant access to a domain user account

1. In the left-hand navigation panel, select the **Users & Groups** page.
2. In the upper right-hand corner, click **Grant Access**.
3. In the **everRun - Grant Access Wizard**, specify the search range in the **Search for** menu.
4. Type the name or group for which to search. Partial names and text are allowed.
5. Click **Search**.
6. Click the green plus sign (+) next to the users or groups you want to add as everRun Availability Console Global Users or Groups of the system.
7. Use the drop-down menus in the Role column to assign a role to the user or group to which you have just granted access. You can assign the following roles:
 - **Administrator** –Enables performance of the full range of system administration activities.
 - **Platform Admin**–Enables Administrator privileges, except for managing user accounts.
 - **Read Only**–Enables read access but no management functions.
 - **Everyone** –Enables only limited read access to certain information.
8. Click **Finish**. The new domain users are displayed in the Grant Access wizard.

To remove access for a domain user account

1. In the left-hand navigation panel, select the **Users & Groups** page.
2. In the upper right-hand corner, click **Grant Access**.
3. In the **everRun - Grant Access Wizard**, click the check box next to users or groups you want to remove.
4. Click **Deny Access**, then click **Finish**.

Related Topic

["Configuring Active Directory" on page 77](#)

4

Chapter 4: Upgrading everRun Software

To upgrade everRun software, use one of the following procedures:



Caution: If your system has external storage, you must call Stratus Customer Assistance to perform an upgrade to everRun Release 7.4.x

- Use an upgrade kit. See ["Upgrading a System Using an Upgrade Kit" on page 115](#).
- Use a DVD. See ["Upgrading a System To Release 7.4.x with a DVD" on page 117](#).

Related Topics

["The Upgrade Kits Page" on page 109](#)

["The everRun Availability Console" on page 58](#)

["Using the everRun Availability Console" on page 57](#)

["Upgrading a System with External Storage from Release 7.3.4 to 7.4.x" on page 284](#)

Upgrading a System Using an Upgrade Kit

This topic describes how to upload an upgrade kit and then use the kit to upgrade everRun software. (If you want to upgrade a system with a DVD, see ["Upgrading a System To Release 7.4.x with a DVD" on page 117](#).)

Prerequisites:



- All PMs and VMs must be in good health before performing an upgrade of everRun software. Before starting an upgrade, examine the everRun Availability Console to verify that there are no alerts indicating PM or VM problems.
- To verify that your system meets the requirements of the upgrade kit, use the **Qualify** button, as described in the following procedure.
- Upgrading from Release 7.3.4 to 7.4.x requires up to 10GiB of additional disk space for the LVM physical volume that is on the boot disk. If you are upgrading a system with external storage from Release 7.3.4 to Release 7.4.x, you must read "[Upgrading a System with External Storage from Release 7.3.4 to 7.4.x](#)" on page 284 before upgrading your system.

Notes:



1. The everRun software upgrade also upgrades the AVCLI software on the everRun system; however, 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 **Downloads** 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 308.
2. Upgrading from Release 7.3.4 to 7.4.x may take approximately two hours depending on the complexity of your configuration and the number of VMs to be migrated during the upgrade. You cannot use the everRun Availability Console while the upgrade is in progress, but the everRun Availability Console may display informational or error messages as needed.

To upgrade the system using 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 - Kit Upload Wizard**.
3. In the **everRun - 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 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 the one to use.
7. Click **Qualify** to verify that your system meets the requirements of the upgrade kit. (This step is recommended, but not required.) After a minute or two, a message appears stating whether or not the kit qualified successfully. If the qualification succeeds, continue with the next step.

If the qualification fails, log on to the host operating system of each PM and see the `/var/opt/ft/log/unity_upgrade.log` file to determine the cause. For example, if there is not enough disk space to complete the upgrade, a message will indicate `Insufficient free space` and report the amount of space needed. If you need help resolving a qualification issue, search for the qualification error message in the **Knowledge Base** in the **Stratus Customer Service Portal** at <https://support.stratus.com>.

8. 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.

Related Topics

["The Upgrade Kits Page" on page 109](#)

["AVCLI Command Descriptions" on page 320 \(see *Kit Information*\)](#)

["The everRun Availability Console" on page 58](#)

["Using the everRun Availability Console" on page 57](#)

["Upgrading a System with External Storage from Release 7.3.4 to 7.4.x" on page 284](#)

Upgrading a System To Release 7.4.x with a DVD

You can upgrade a system, with or without external storage, from everRun Release 7.3.4.0 to Release 7.4.x using a DVD.



Caution: If your system has external storage, you must call Stratus Customer Assistance to perform an upgrade to everRun Release 7.4.x.



Prerequisites:

- All PMs and VMs in must be in good health before performing an upgrade of everRun software. Before starting an upgrade, examine the everRun Availability Console to verify that there are no alerts indicating PM or VM problems.
- Upgrading from Release 7.3.4 to 7.4.x requires up to 10GiB of additional disk space for the LVM physical volume that is on the boot disk. Before you begin the upgrade, make sure that your system has at least 10GiB of additional disk space.



Note: Upgrading from Release 7.3.4 to 7.4.x may take approximately two hours depending on the complexity of your configuration. You cannot use the everRun Availability Console while the upgrade is in progress, but the everRun Availability Console may display informational or error messages as part of the procedure.

The procedure includes shutting down VMs and then sequentially on each node: putting the node in maintenance mode, shutting down the node, powering on the node and using the BIOS setup utility to boot from DVD and then recover the node. Note that you cannot use the UEFI setup utility to upgrade from a DVD.

For detailed information, access the Knowledge Base (see "[Accessing Knowledge Base Articles](#)" on page 481) to search for the article *Upgrading a System from everRun Release 7.3.4.0 to 7.4.x.x With a DVD* (KB-4346).

Related Topics

["Upgrading everRun Software" on page 115](#)

["The everRun Availability Console" on page 58](#)

["Using the everRun Availability Console" on page 57](#)

["Upgrading a System with External Storage from Release 7.3.4 to 7.4.x" on page 284](#)

5

Chapter 5: Managing Logical Disks

Manage logical disks using the everRun Availability Console. For overview information see ["Logical Disk Management" on page 119](#) 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 120](#)
- ["Activating a New Logical Disk" on page 122](#)
- ["Creating a New Storage Group" on page 123](#)
- ["Deleting a Storage Group" on page 123](#)
- ["Assigning a Logical Disk to a Storage Group" on page 124](#)

For information on working with external storage, see ["Managing External Storage " on page 279](#).

Logical Disk Management

In an everRun 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 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 122](#).

You need to respond to alerts regarding missing or failed logical disks. everRun software may detect a logical disk failure when a physical disk is removed or fails. The everRun 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 95](#).

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 120](#)

Related Topics

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

["The everRun Availability Console" on page 58](#)

["Storage Requirements" on page 26](#)

Responding to a Failed Logical Disk

When everRun 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 119](#).) 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. Systems configured for UEFI will only boot from the logical disk that the everRun software was originally installed on.
4. In some legacy BIOS 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 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 122](#).

Related Topics

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

["The everRun Availability Console" on page 58](#)

Activating a New Logical Disk

In an everRun 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.



Prerequisite: Note any relevant requirements in ["Storage Requirements" on page 26](#).

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 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 and makes it available to be added to a storage group.

Related Topics

["Responding to a Failed Logical Disk" on page 120](#)

["Logical Disk Management" on page 119](#)

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

["The everRun Availability Console" on page 58](#)

["Storage Requirements" on page 26](#)

Creating a New Storage Group

You can create a new storage group. When you create a new storage group, it will have no logical disks assigned to it. See ["Assigning a Logical Disk to a Storage Group" on page 124](#) for more information.

To create a new storage group

1. On the **STORAGE GROUPS** page, click the **Create** button. The **New Storage Group** dialog box appears.
2. In the **Name** box, type the name of the new storage group.
3. Click **Create Storage Group**.

Related Topics

["Assigning a Logical Disk to a Storage Group" on page 124](#)

["Deleting a Storage Group" on page 123](#)

["Storage Requirements" on page 26](#)

Deleting a Storage Group

You can delete a storage group as long as no logical disks are assigned to it.



Note: If you attempt to delete a storage group that has one or more logical disks assigned to it, the system will prompt you to first move the logical disks to another storage group and then perform the Delete operation.

To delete a storage group

1. On the **Storage Groups** page, select the storage group you wish to delete.
2. Click the **Delete** button. The **Confirm** dialog box appears.
3. Click **Yes** to delete the storage group.

Related Topics

["Creating a New Storage Group" on page 123](#)

Assigning a Logical Disk to a Storage Group

You can assign an empty logical disk to a storage group.

Notes:



- The logical disk that you want to add must be empty. Its **Size Used** value must be zero.
- The sector size (4K or 512B) of the logical disk must match the sector size of the storage group.
- Note any relevant requirements in ["Storage Requirements" on page 26](#).

To assign a logical disk to a storage group

1. On the **PHYSICAL MACHINES** page, select **node 0**.
2. Click the **Storage** tab.
3. Select an empty logical disk (its **Size Used** is 0).
4. In the **Actions** column, click **Move to Storage Group**.
5. **The Move Disk to Storage Group** dialog box appears. In the **Storage Group** drop down box, select the storage group.
6. Click **Move to Storage Group**.
7. On the **PHYSICAL MACHINES** page, select **node 1**.
8. Repeat steps 2 through 6 above.

On the **STORAGE GROUPS** page, the new storage group will appear with a non-zero size.

Related Topics

- ["Deleting a Storage Group" on page 123](#)
- ["Storage Requirements" on page 26](#)

6

Chapter 6: 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 95](#).

Many of the tasks you perform from the **Physical Machines** page require maintenance mode; for information, see ["Maintenance Mode" on page 125](#).

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

- ["Rebooting a Physical Machine" on page 127](#)
- ["Shutting Down a Physical Machine" on page 127](#)
- ["Load Balancing" on page 129](#)

To troubleshoot a PM by recovering a failed PM or resetting MTBF for a failed machine, see ["Troubleshooting Physical Machines" on page 129](#).

To perform maintenance tasks, see ["Maintaining Physical Machines" on page 269](#).

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 VMs running on it to the other PM, which protects the VMs from any potential disruption caused by the service.

When the primary PM (**nodex (primary)**) enters maintenance mode, the other PM becomes primary.

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 127.](#))



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 58](#)

["Managing Physical Machines" on page 125](#)

["Physical Machines and Virtual Machines" on page 9](#)

["The Physical Machines Page" on page 95](#)

["The Virtual Machines Page" on page 99](#)

Rebooting a Physical Machine

Reboot a physical machine (PM) to restart its everRun 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 61.](#))

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 125](#)

["The everRun Availability Console" on page 58](#)

["Managing Physical Machines" on page 125](#)

["The Physical Machines Page" on page 95](#)

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 62](#).
2. Do not use the `-f` (force) option with the `halt`, `poweroff`, or `reboot` command of the CentOS operating system. Doing so causes FT guests that are active on the same node to hang. Instead, use the everRun Availability Console and maintenance mode to shutdown a PM, as described in the procedure below.
3. The everRun 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. **Manually power off a PM 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 125](#)

["The everRun Availability Console" on page 58](#)

["Managing Physical Machines" on page 125](#)

["The Physical Machines Page" on page 95](#)

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

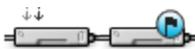
Click **Load Balance** to initiate automatic load balancing of a VM.

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 265](#)

Troubleshooting Physical Machines

The following topics describe troubleshooting procedures for PMs:

- ["Recovering a Failed Physical Machine" on page 130](#)
- ["Recovering a Failed Physical Machine with External Storage " on page 132](#)
- ["Resetting MTBF for a Failed Physical Machine" on page 132](#)

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 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 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 re-creates the /boot and root file systems, re-installs CentOS and the everRun software, and attempts to connect to an existing system.



Warning: This procedure deletes all software you may have installed on the PM and all PM configuration information you may have entered before recovery. After you complete this procedure, you must manually re-install all your software and reconfigure the PM to match your original settings.



Note: If you need to repair or replace the PM, see ["Replacing Physical Machines, Motherboards, NICs, or RAID Controllers" on page 274](#), which requires the **Remove** operation of a node in maintenance mode.

Prerequisites:

1. Determine which PM you need to recover.
2. Obtain installation software for the everRun 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_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 Software" on page 36](#).

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 firmware (BIOS or UEFI) setup utility 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.
8. Manually reinstall applications and any other host-level software and reconfigure the PM to match your original settings.

Related Topics

["Maintenance Mode" on page 125](#)

["Managing Physical Machines" on page 125](#)

["The everRun Availability Console" on page 58](#)

["The Physical Machines Page" on page 95](#)

Recovering a Failed Physical Machine with External Storage

You can recover a failed physical machine (PM) that is connected to external storage while the system is running. For systems with Fibre Channel or iSCSI external storage, you begin by recovering the PM, and then you update `initrd` and modify the `multipath.conf` and `lvm.conf` files. For systems with iSCSI external storage, you also stop and restart the `everrun-sim` service, remap LUNs, and use `iscsiadmin` for various tasks.

For detailed information, access the Knowledge Base to search for the article *Recovering a Failed Physical Machine with External Storage* (KB-4300). See ["Accessing Knowledge Base Articles" on page 481](#).



Note: If you need to repair or replace a PM that is attached to external storage, see ["Replacing a Failed Physical Machine Connected to External Storage" on page 277](#).

Related Topics

["Recovering a Failed Physical Machine" on page 130](#)

["External Storage" on page 17](#)

["Managing External Storage " on page 279](#)

["Managing Physical Machines" on page 125](#)

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 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 125](#)

["Managing Physical Machines" on page 125](#)

["The everRun Availability Console" on page 58](#)

["The Physical Machines Page" on page 95](#)

7

Chapter 7: 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 99](#). To perform specific management tasks, see the following topics.

To manage the operational state of a VM, see:

- ["Starting a Virtual Machine" on page 225](#)
- ["Shutting Down a Virtual Machine" on page 225](#)
- ["Powering Off a Virtual Machine" on page 226](#)
- ["Opening a Virtual Machine Console Session" on page 227](#)
- ["Renaming a Virtual Machine" on page 228](#)
- ["Removing a Virtual Machine" on page 229](#)

To display information about a VM, use the `snmptable` command (see ["Obtaining System Information with snmptable" on page 484](#)).

To create or configure a VM, see:

- ["Planning Virtual Machine Resources" on page 136](#) (virtual CPUs, memory, storage, and networks)
- ["Creating and Migrating Virtual Machines" on page 141](#)
- ["Managing Snapshots" on page 249](#)

- ["Managing Virtual CDs" on page 244](#)
- ["Configuring Windows-based Virtual Machines" on page 215](#)
- ["Configuring Linux-based Virtual Machines" on page 221](#)
- ["Managing Virtual Machine Resources" on page 230](#)

To complete advanced tasks, see:

- ["Assigning a Specific MAC Address to a Virtual Machine" on page 263](#)
- ["Selecting a Preferred PM for a Virtual Machine" on page 265](#)
- ["Changing the Protection Level for a Virtual Machine \(HA or FT\)" on page 265](#)
- ["Configuring the Boot Sequence for Virtual Machines" on page 266](#)
- ["Resetting MTBF for a Failed Virtual Machine" on page 267](#)

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 136](#)
- ["Planning Virtual Machine Memory" on page 138](#)
- ["Planning Virtual Machine Storage" on page 139](#)
- ["Planning Virtual Machine Networks" on page 141](#)

Planning Virtual Machine vCPUs

Allocate virtual CPUs (vCPUs) to assign computing resources to a virtual machine (VM) on your everRun 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 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 all VMs is equal to the total number of vCPUs available on the everRun system minus the number of vCPUs allocated to the everRun system software (you can set system vCPUs to 2 or 4, as described in ["Configuring System Resources" on page 76](#)). 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 available to all VMs minus the number of vCPUs allocated to currently running VMs (as described above) within the limits listed in ["Virtual Machine Recommendations and Limits" on page 478](#).
- 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 231](#)), 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 61](#)) 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 24](#)

["Creating and Migrating Virtual Machines" on page 141](#)

["Managing Virtual Machine Resources" on page 230](#)

Planning Virtual Machine Memory

Allocate memory to assign physical memory to a virtual machine (VM) on your everRun 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 27](#)) 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 76](#)). 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 61](#)) 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 27](#)

["Creating and Migrating Virtual Machines" on page 141](#)

["Managing Virtual Machine Resources" on page 230](#)

Planning Virtual Machine Storage

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

When you configure the everRun system, you create storage groups from the available logical disks. You allocate VM volumes and virtual CDs (VCDs) from these storage groups, 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. 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).

- Observe 512B and 4K sector size restrictions

Stratus recommends using disks with the 4K sector size for better performance. When using 4K storage, ensure that your storage groups support the sector size of the VM volumes that you want to create or import:

- A storage group with the 512B sector size supports only 512B VM volumes.
- A storage group with the 4K sector size supports 4K or 512B VM volumes.

Note that the boot volume for each VM must be 512B, regardless of the storage group's sector size. Only data volumes can use the 4K sector size. Ensure that your guest operating systems support 4K volumes before creating or attaching them.

- Leave storage 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.)

- Leave storage space for VM snapshots

When you create each VM volume, you specify its volume size as well as the size of the larger volume container in which the volume and its snapshots are stored. To leave enough space for the snapshots that you plan to create, start by allocating a volume container that is at least two times the size of the volume that it contains; however, your needs may vary depending on VM snapshot activity. For more information about estimating the amount of storage needed in a volume container, see ["Sizing Volume Containers" on page 16](#).

To conserve storage space in a volume container, you can remove older or obsolete snapshots as described in ["Removing a Snapshot" on page 262](#). If necessary, you can also expand a volume container as described in ["Expanding a Volume Container on the everRun System" on page 241](#).

- 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 26](#)

["Creating and Migrating Virtual Machines" on page 141](#)

["Managing Virtual Machine Resources" on page 230](#)

Planning Virtual Machine Networks

Plan network resources to determine how you will allocate available virtual networks to the virtual machines (VMs) on your everRun 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 27](#)

["Creating and Migrating Virtual Machines" on page 141](#)

["Managing Virtual Machine Resources" on page 230](#)

Creating and Migrating Virtual Machines

Generate a new virtual machine (VM) on the everRun 7.x 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 or Avance VM.

To create a new VM (without an existing source VM or PM), see ["Creating a New Virtual Machine" on page 143](#).

To copy an existing VM on your system for the purpose of creating a new VM or creating a duplicate VM for troubleshooting, see ["Copying a Virtual Machine" on page 149](#).

To migrate or import a system from a non-everRun 7.x source, see ["Migrating From Avance or everRun MX Systems" on page 160](#) for considerations, and then see one of the following topics depending on your needs:

- ["Migrating a Physical Machine or Virtual Machine to an everRun 7.x System" on page 151](#) (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 system.
- ["Importing an OVF File from an everRun MX System" on page 182](#) (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 7.x system.
- ["Importing an OVF File from an Avance System" on page 190](#) (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 7.x system.

To migrate or import a VM from another everRun 7.x system, or to restore a VM on the same everRun 7.x system, see one of the following topics:

- ["Migrating a Physical Machine or Virtual Machine to an everRun 7.x System" on page 151](#)
Use the *P2V client (virt-p2v)* to transfer a VM directly over the network to a new VM on another everRun 7.x system or on the same everRun 7.x system.
- ["Exporting a Virtual Machine from an everRun System" on page 208](#)

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

- ["Managing Snapshots" on page 249](#)

Use the everRun Availability Console to create a snapshot of the source VM, then use the snapshot to create a new VM on the same everRun system or export the snapshot to OVF and VHD files on a network share.

- ["Importing an OVF File from an everRun 7.x System " on page 198](#)

Use the everRun Availability Console to import OVF and VHD files to another everRun 7.x system or back to the same everRun 7.x system.

["Replacing a Virtual Machine from an OVF File" on page 203](#)

Use the everRun Availability Console to import OVF and VHD files back to the same everRun 7.x system to overwrite and restore an existing VM from a previous backup copy.

Related Topics

["Managing Virtual Machines" on page 135](#)

Creating a New Virtual Machine

Create a new virtual machine (VM) to install a guest operating system on your everRun system. (You can also migrate an existing VM or physical machine (PM), as summarized in ["Creating and Migrating Virtual Machines" on page 141.](#))

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.

Note:



If you need to create a Windows Server 2003 VM, see ["Creating a New Windows Server 2003 Virtual Machine" on page 147](#). You must perform a different procedure to create a Windows Server 2003 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 136](#).
- You can select a remote ISO or a bootable virtual CD (VCD) as the source that the VM boots from. For a remote ISO, you must have a URL or path name for the repository; and for a remote ISO on a shared network drive, you must have a user name and password. If you need a bootable VCD of the Windows or Linux installation media, create it as described in ["Creating a Virtual CD" on page 244](#). The bootable VCD must be a single CD or DVD. Multiple CDs or DVDs are not supported.
- Ensure that both PMs of the everRun 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 95](#)), 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 99](#)), 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:
 - **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.
 - **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.

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

- c. For **Boot From**, select one of the following as the boot source:
 - **VCD**—The boot source is a VCD. Select a source from the pull-down menu.
 - **Remote ISO via Windows Share (CIFS/SMB)**—The boot source is a remote ISO file on a shared network drive. You must enter values for **User Name** and **Password**. For **Repository**, enter a value in the format *//iso_url/path_to_ISO_folder*.
 - **Remote ISO via NFS**—The boot source is a remote ISO file, accessed through NFS. For **Repository**, enter a value in the format *//iso_url/path_to_ISO_folder*.

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

- d. Click **Next**.
4. On the **vCPUs and Memory** page:
 - a. Specify the number of **vCPUs** and the amount of **Memory** to assign to the VM. For more information, see ["Planning Virtual Machine vCPUs" on page 136](#) and ["Planning Virtual Machine Memory" on page 138](#).
 - b. If applicable, select the **Keyboard Layout in Console** to set the keyboard type for your country.
 - c. Click **Next**.
 5. On the **Volumes** page:
 - a. Type the **Name** of the boot volume as it will appear in the everRun Availability Console.
 - b. Type the **Container Size** and **Volume Size** of the volume to create in gigabytes (GB). The container size is the total size for the volume including extra space to store snapshots. The volume size is the portion of the container that is available to the guest operating system. For more information about allocating storage, see ["Sizing Volume Containers" on page 16](#) and ["Planning Virtual Machine Storage" on page 139](#).

- c. Select the **Disk Image** format:
 - **RAW**—raw disk format
 - **QCOW2**—QEMU Copy On Write (QCOW2) format, which supports snapshots
- d. Select the **Storage Group** in which to create the volume, and, if applicable, select the **Sector Size**.

Ensure that you select a storage group that supports the sector size of the volume that you want to create (see ["Planning Virtual Machine Storage" on page 139](#)). Note that the boot volume must have a sector size of 512 B. You can select the sector size, either 4K or 512B, only for data disks.

- e. If applicable, create additional data volumes by clicking **Add New Volume** and specifying the parameters for each volume. (You can also add volumes after you create the VM by using the **Reprovision Virtual Machine** wizard, as described in ["Creating a Volume in a Virtual Machine" on page 234](#).)
 - f. Click **Next**.
6. On the **Networks** page, select the shared networks to attach to the VM (for more information, see ["Planning Virtual Machine Networks" on page 141](#)). 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 (allow pop-ups in your browser, if necessary) and respond to any prompts in the VM console session.
 9. After you install the operating system, configure the additional resources and software necessary for production use, as described in:
 - ["Configuring Windows-based Virtual Machines" on page 215](#)
 - ["Configuring Linux-based Virtual Machines" on page 221](#)



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

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

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

Related Topics

["Copying a Virtual Machine" on page 149](#)

["Renaming a Virtual Machine" on page 228](#)

["Removing a Virtual Machine" on page 229](#)

["Creating and Migrating Virtual Machines" on page 141](#)

["Managing Virtual Machine Resources" on page 230](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

Creating a New Windows Server 2003 Virtual Machine

Follow these instructions to create a new Windows Server 2003 VM on your everRun system. You should understand the following considerations before you create the Windows Server 2003 VM:

- The Windows Server 2003 operating system is no longer supported by Microsoft.
- The **only version** of Windows Server 2003 that everRun systems support is the **Windows Server 2003 R2 Enterprise SP2 32-bit** operating system.
- The network VirtIO driver is not automatically installed as it is when creating VMs with other operating systems. The manual steps required to do this are provided in the procedure below.



Note: The following procedure documents only the unique actions that are required to install this guest OS on your everRun system. You must also respond to other normal installation prompts not documented here (for example choosing a locale).

To create a new Windows Server 2003 VM

1. Create a bootable virtual CD (VCD) of the Windows Server 2003 media as described in ["Creating a Virtual CD" on page 244](#).

2. Perform steps 1 through 7 of the procedure as described in ["Creating a New Virtual Machine" on page 143](#).
3. When the dialog box informing you that the software has not passed Windows Logo testing appears click **Yes** to continue the installation.
4. When the dialog box informing you that the RedHat VirtIO SCSI controller driver has not passed Windows Logo testing appears, click **Yes** to continue the installation.
5. When the dialog box informing you that Windows Setup is not complete, click **Cancel**.
6. In the Windows Setup message informing you that you have chosen to not continue Setup, click **OK**.
7. Open **Computer Management** and click **Device Manager**.
8. In the right-hand pane of Computer Manager, under **Other Devices**, right click **Ethernet Controller**. In the pop-up menu click **Update Driver**.
9. In the **Hardware Update Wizard**, select **No, not this time**. Click **Next**.
10. In the **Hardware Update Wizard**, select **Install from a list or specific location (Advanced)**. Click **Next**.
11. In the **Hardware Update Wizard**, select **Search removable media (floppy, CD-ROM...)**. Click **Next**.
12. In the **Hardware Update Wizard**, select the top-most Red Hat VirtIO Ethernet Adapter entry. Click **Next**.
13. When the Hardware Installation message informs you that the software has not passed Windows Logo testing click **Continue Anyway**, then click **Finish**.
14. Close **Computer Management**.
15. Shut down the VM that was just installed. This must be done in order to remove the virtual floppy drive that was automatically installed earlier during the installation.



Note: If you need to install optional software from the Windows Server CD2 you must obtain an ISO image of that media. Place that ISO image on a network that the system can access, and run the setup.exe file.

16. 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 215](#).

Related Topics

["Creating and Migrating Virtual Machines" on page 141](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

Copying a Virtual Machine

Copy a virtual machine (VM) if you want to clone an existing VM on your everRun system. For example, you can copy a healthy VM to create a new VM, or you can copy a VM that is not working properly and use the copy for troubleshooting purposes. (If you want to import or migrate a VM from a different system, see the overview in ["Creating and Migrating Virtual Machines" on page 141](#).)

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

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

Notes:



- Copying a VM does not copy the snapshots from the source VM, but you can configure the container size of the new VM to make it possible to create new snapshots.
- To prevent conflicts with the source VM, the copy wizard automatically assigns a new MAC address to each network interface in the new VM; however, you may need to manually update any IP addresses and host names.
- If the everRun system switches from the primary PM to the secondary PM while copying a VM, the copy process fails. This does not affect the continuous uptime of your system, but you must delete any volumes associated with the copied VM and start the copy again.

Prerequisites:



- You must shut down a VM before copying it.
- Both PMs of the everRun system must be online for the copy process to function properly.

To copy a VM on the everRun system

1. On the **Physical Machines** page (see ["The Physical Machines Page" on page 95](#)), 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 99](#)), select the VM that you want to copy and click **Shutdown**.
3. When the VM has stopped, click **Copy** to open the copy wizard.
4. On the **Name, Description, and Protection** page:
 - a. Type the **Name** and an optional **Description** for the VM as they will appear in the everRun Availability Console.
 - b. Select the level of protection to use for the VM:
 - **Fault Tolerant (FT)**
 - **High Availability (HA)**

For information about these levels of protection, see ["Creating a New Virtual Machine" on page 143](#) and ["Modes of Operation" on page 11](#).
 - c. Click **Next**.
5. On the **vCPUs, Memory, and Console Keyboard Layout** page:
 - a. Specify the number of **vCPUs** and the amount of **Memory** to assign to the VM. For more information, see ["Planning Virtual Machine vCPUs" on page 136](#) and ["Planning Virtual Machine Memory" on page 138](#).
 - b. If applicable, select the **Keyboard Layout in Console** to set the keyboard type for your country.
 - c. Click **Next**.
6. On the **Volumes** page, you can:
 - Type the volume **Name**.
 - Specify the **Container Size** and **Volume Size** of each volume to make room for snapshots.
 - Assign each volume to a **Storage Group**.
 - Specify a **Sector Size**.

- Click **Add New Volume** to create a new data volume. (If the button is not visible, scroll down to the bottom of the wizard page.)

For more information, see ["Planning Virtual Machine Storage" on page 139](#). To continue, click **Next**.

7. On the **Networks** page, activate the check box for each shared network that you want to attach to the VM.
8. On the **Copy Summary** page:
 - a. Review the configuration summary. If you need to make changes, click **Back**.
 - b. To proceed with copying the VM, click **Finish**.

After the copy process is complete; the everRun 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 copy process.

To clean up after a canceled or failed copy process

Remove any volumes associated with the copied VM.

Related Topics

["Creating and Migrating Virtual Machines" on page 141](#)

["Managing Virtual Machine Resources" on page 230](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

Migrating a Physical Machine or Virtual Machine to an everRun 7.x System

Migrate a physical machine (PM) or virtual machine (VM) to transfer it over the network to a new VM on the everRun 7.x system. (You can also import an Open Virtualization Format (OVF) file to the everRun 7.x system, as summarized in ["Creating and Migrating Virtual Machines" on page 141](#).)

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 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 and 7, Windows 7 and 8.x, and Windows Server 2008 R2 and 2012. (If you need to migrate a Windows Server 2003 VM, see ["Migrating a Windows Server 2003 VM to an everRun 7.x System" on page 179](#). You must perform a different procedure to migrate a Windows Server 2003 VM.)
- For Windows-based PMs or VMs that support *fast startup* mode, you must disable this feature before the migration process; otherwise, the migration process will fail as described below in **Troubleshooting**.
- For Linux-based PMs or VMs, consider editing the `/etc/fstab` file before the migration process to comment out entries for data volumes and allow only the boot volume to mount. Because Linux-based VMs use different device names on the everRun system, your new VM may boot into single-user mode if it cannot the 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**.
- Your source PM or VM must be offline for the duration of the migration process. Consider scheduling a planned maintenance period for the migration.
- The time required for the 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 system. For example, transferring a source system with one 20 GB boot volume 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.
- To prevent conflicts with the original PM or VM, the P2V client automatically assigns a new MAC address to each network interface in the new VM; however, you must manually update any IP addresses and host names as needed.
- If the everRun 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 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 system

1. Download the P2V client ISO file from the **Downloads** page at <http://www.stratus.com/go/support/everrun>.
 - a. On the **Downloads** page, click **everRun** (if it is not already displayed) and then select the appropriate version.
 - b. Scroll down to **Drivers and Tools** and then continue scrolling to **everRun P2V Client for Virtual or Physical Machine Migration**.
 - c. Select the **P2V Client (virt-p2v)** file.
2. If you want to verify the integrity of the ISO image, also download the associated `fciv` checksum file, and then download the Microsoft File Checksum Integrity Verifier (FCIV) executable file from the Microsoft Support web site. Save both files to the directory that contains the downloaded ISO file.

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

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

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

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

To prepare for migrating a VM to the everRun system

1. Download the P2V client ISO file from the **Drivers and Tools** section of the **Downloads** page at <http://www.stratus.com/go/support/everrun>. Ensure that you download the version of the P2V client that matches the version of the everRun system to which you are migrating

the VM.

2. If you want to verify the integrity of the ISO image, also download the associated `fciv` checksum file, and then download the Microsoft File Checksum Integrity Verifier (FCIV) executable file from the Microsoft Support web site. Save both files to the directory that contains the downloaded ISO file.

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

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

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

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

To migrate a PM or VM to the everRun system

1. Power on the source PM or VM to boot the P2V client. After a minute or so, the **virt-p2v** window is displayed.
2. The P2V client automatically obtains network settings through DHCP. Static settings are unnecessary for the migration process, but you can optionally click **Configure network** to specify the settings. (If necessary, configure the network settings of the target VM later on the everRun system.)
3. Enter the connection settings for the **Conversion server** (the everRun 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 21](#).)
4. Click **Test connection**. If the P2V client connects to the everRun system, click **Next** to continue. The **Target properties** page is displayed.

If the P2V client cannot connect, verify the connection settings and try to connect again.
5. 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 system.)

6. The **# vCPUs** and **Memory(MB)** values are automatically detected and completed, but optionally modify them if you want the VM on the everRun system to have more CPUs or memory than the source PM or VM.
7. Specify the **Virt-v2v output options** for the target VM, as follows:
 - a. Next to **Output to**, select **everrunha** (High Availability) or **everrunft** (Fault Tolerant) operation. (For information about these levels of protection, see ["Creating a New Virtual Machine" on page 143](#) and ["Modes of Operation" on page 11.](#))
 - b. Next to **Output format**, select the disk image format, **raw** or **qcow2**. (The qcow2 format supports snapshots.)
8. If you want to save debugging messages from the migration process, optionally select the **Enable server-side debugging** check box. (The debugging messages are included if you generate a diagnostic file for your authorized Stratus service representative, as described in ["Creating a Diagnostic File" on page 85.](#))
9. Select which **Fixed hard disks** (volumes) to include in the migration by activating the check box next to each device.

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

If the target everRun system has more than one storage group, you can also select the storage group in which to create each volume. Double-click the volume entry to open the **Choose Storage Group** panel. Ensure that you select a **Storage Group** that supports the sector size of the volume you are importing (see ["Planning Virtual Machine Storage" on page 139](#)) and select the **Sector Size** that matches the source volume (the P2V client cannot convert the sector size of a volume). Note that the boot volume must have a sector size of 512 B. You can select the sector size, either 4K or 512B, only for data disks.

10. Select which **Network Interfaces** to include in the migration by activating the check box next to each device.

If the target everRun system has more than one shared network, you can also select the shared network to connect with each network interface. Double-click the network interface to open the **Choose Shared Network** panel, select the shared network, and click **OK**.

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



Note: After the migration, the new VM on the everRun 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 system

1. Open the **Virtual Machines** page (see ["The Virtual Machines Page" on page 99](#)) 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 231](#). 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 227](#).)
5. For Windows-based VMs only, update the VirtIO drivers to the latest supported versions, as described in ["Updating the VirtIO Drivers \(Windows-based VMs\)" on page 215](#). (The correct VirtIO drivers are already present in Linux-based VMs.)



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

6. Disable any guest operating system services that are unnecessary for operation on the everRun system:
 - If you migrated from a PM source, disable any services that interact directly with hardware. Examples include:
 - Dell OpenManage (OMSA)
 - HP Insight Manager
 - Diskkeeper
 - If you migrated from a VM source, disable any services associated with other hypervisors. Examples include:
 - VMware Tools
 - Hyper-V Tools
 - Citrix Tools for Virtual Machines

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

7. If necessary, update the network configuration settings in the guest operating system and restart it to enable the settings.
8. Verify that you have configured your guest operating system with the additional Windows- or Linux-based system settings described in:
 - ["Configuring Windows-based Virtual Machines" on page 215](#)
 - ["Configuring Linux-based Virtual Machines" on page 221](#)

After you verify that the new VM is functioning properly, the migration process is complete; however, the everRun 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 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 85](#). The diagnostic file includes any server-side debugging messages from the migration process.

To recover from a migration that fails with the error message, Failed to mount

'/dev/sda1: Operation not permitted

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

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

To resolve the issue, log on to the guest operating system of the source VM, disable fast startup mode, and shut down the operating system. Then, restart the migration process.

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

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

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

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

To recover missing data volumes in the VM on the everRun system

If your data volumes do not appear in the VM on the everRun 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 (`/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 system

If your network devices do not appear in the VM on the everRun 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 Avance or everRun MX Systems" on page 160](#)

["Creating and Migrating Virtual Machines" on page 141](#)

["Configuring Windows-based Virtual Machines" on page 215](#)

["Configuring Linux-based Virtual Machines" on page 221](#)

["Managing Virtual Machine Resources" on page 230](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

Migrating From Avance or everRun MX Systems

If you are migrating from an everRun MX system or an Avance unit to an everRun 7.x system, and you want to transfer virtual machines (VMs) from the other system, see ["Creating and Migrating Virtual Machines" on page 141](#).

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

- ["Planning to Migrate from an everRun MX System" on page 161](#) (System-to-System Migration)
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 7.x system.
- ["Converting an everRun MX System to an everRun 7.x System" on page 164](#) (In-Place Migration)
Use this procedure to perform an in-place migration of an everRun MX system and its VMs to the everRun 7.x software.
- ["Planning to Migrate from an Avance Unit" on page 171](#) (System-to-System Migration)
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 7.x system.
- ["Converting an Avance Unit to an everRun 7.x System" on page 173](#) (In-Place Migration)
Use this procedure to perform an in-place migration of an Avance unit and its VMs to the everRun 7.x software.

Related Topics

["Planning" on page 23](#)

["Software Installation" on page 34](#)

["Post-Installation Tasks" on page 53](#)

["Creating and Migrating Virtual Machines" on page 141](#)

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 7.x system.

For all systems, see ["Creating and Migrating Virtual Machines" on page 141](#) for information about migrating your virtual machines (VMs) to the everRun 7.x system.



Note: For best results, contact your authorized Stratus service representative for assistance in evaluating and performing upgrades from an everRun MX 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 7.x systems, as described in ["Physical Machine System Requirements" on page 476](#).

Multi-node XenServer pools are supported by everRun MX, but only two-node configurations are supported for everRun 7.x 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 7.x software. See ["Compatible Guest Operating Systems" on page 474](#).

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 an everRun 7.x System" on page 151](#)) or the import process (as described in ["Importing an OVF File from an everRun MX System" on page 182](#)).

Network Preparation

Prepare the platform network and the networking environment to conform to the everRun 7.x requirements. See ["General Network Requirements and Configurations" on page 27](#).

Management Network Access

The XenServer management network becomes the everRun 7.x 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 7.x 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 7.x 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 7.x nodes as necessary by the everRun 7.x 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 7.x 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 7.x 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 private network must be identified. Only one everRun 7.x 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 7.x nodes.

On the everRun 7.x 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 redundant-path storage, but this is not supported on an everRun 7.x system.

For physical storage requirements, see ["Storage Requirements" on page 26](#).

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 7.x 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

After the nodes in the everRun 7.x system are configured, you can install and configure the everRun 7.x software, as described in ["Software Installation" on page 34](#).

Migrating Virtual Machines

Using either the P2V client migration process or the OVF import process, migrate the VMs to the everRun 7.x system. For an overview of each process, see ["Creating and Migrating Virtual Machines" on page 141](#).

Converting an everRun MX System to an everRun 7.x System

Convert an everRun MX system to an everRun 7.x system to perform an in-place migration of the everRun MX system and its virtual machines (VMs) to the everRun 7.x software.

To convert an everRun MX system, shut down one physical machine (PM) or *node* in the everRun MX system and install the everRun 7.x software on that node. Use the P2V client to transfer each VM from the everRun MX node to the everRun 7.x node over the network. Then, install the everRun 7.x software on the remaining node.



Caution: Consider backing up the everRun MX system and its VMs and recording its settings before the conversion. Converting an everRun MX system to an everRun 7.x system ultimately overwrites everything on your everRun MX system (after you migrate your VMs to the everRun 7.x node).



Notes:

- For best results, contact your authorized Stratus service representative for assistance in evaluating and performing upgrades from an everRun MX system.
- Before converting an everRun MX system to an everRun 7.x system, verify that your PMs and VMs are supported as described in ["Physical Machine System Requirements" on page 476](#) and ["Compatible Guest Operating Systems" on page 474](#).

To prepare for converting an everRun MX system

1. Plan for converting your everRun MX system by reviewing the information:

- ["Planning to Migrate from an everRun MX System" on page 161](#)

Describes some items to consider when migrating or converting from an everRun MX system to an everRun 7.x system.

- ["Software Installation" on page 34](#)

Summarizes the steps for installing the everRun 7.x software.

- ["Migrating a Physical Machine or Virtual Machine to an everRun 7.x System" on page 151](#)

Describes how to use the P2V client to migrate a VM from one system to another. Also describes some steps you may need to complete in your guest operating systems **before** migrating the VMs to ensure that the VMs will function properly on the everRun 7.x system.

2. Back up your everRun MX system and VMs.
3. Download the everRun 7.x ISO file from the **Downloads** page at <http://www.stratus.com/go/support/everrun>.
 - a. On the **Downloads** page, click **everRun** (if it is not already displayed) and select the appropriate version.
 - b. Scroll down to **Product Downloads**, and under **Install**, click the link to the appropriate ISO image (everRun_install-7.x.x.x-xxx.iso).
 - c. Save the ISO image
4. Download the P2V client ISO file from the same **Downloads** page.
 - a. On the **Downloads** page, click **everRun** (if it is not already displayed) and then select the appropriate version.
 - b. Scroll down to the **Drivers and Tools** bar and then continue scrolling to **everRun P2V Client for Virtual or Physical Machine Migration**.
 - c. Select the **P2V Client (virt-p2v)** file.
 - d. Save the file.
5. Burn the everRun 7.x ISO file to a physical DVD that you will use to install the everRun 7.x software on each PM in your system.

6. Burn the P2V client ISO file to a physical CD that you will boot in each everRun MX VM to transfer the VMs to the everRun 7.x system.
7. Contact your network administrator to request at least one static IP address to use as the system-wide IP address for your converted everRun 7.x system. Request an additional static IP address for each of the two nodes if you do not have a DHCP server to automatically assign these addresses or if you prefer to use only static addresses.



Note: You must maintain unique system IP addresses for the everRun MX system and the everRun 7.x system while both systems are online; however, if you want to re-use the original everRun MX system IP address for the everRun 7.x system, you can change the network settings of the everRun 7.x system after the conversion is complete.

To shut down the master server of the everRun MX system

Starting with both nodes running the everRun MX software, do the following:

1. Log on to the **everRun Availability Center** with the hostname or IP address of your everRun MX master node at:

`http://everRunMX-address:8080`

2. On left-hand navigation panel, click the **Hosts** tab.
3. Right-click the master server and select **Shutdown**.
4. Allow the server to evacuate the VMs and properly shut down. You can watch its progress on the **everRun Log** tab.

When the server is shut down, a message describes that you have lost connection to the everRun Availability Center. This is expected.

5. Open **Citrix XenCenter** and connect to remaining server of the everRun MX system, which is now the master server.
6. Make sure the VMs are still running on the remaining server before continuing.

To convert the first node of the everRun MX system to an everRun 7.x node



Caution: Converting a node to the everRun 7.x software erases all hard drives in that node.

Starting one node shut down and the second node running the everRun MX software, do the following:

1. Insert the everRun 7.x DVD into the physical DVD drive of the offline node and boot the node to start the installation program.
2. Follow the instructions in ["Installing Software on the First PM" on page 45](#) to install the everRun 7.x software on the first node. Power on the node, update the necessary settings in the firmware (BIOS or UEFI) setup utility, and boot the node from the everRun 7.x DVD to run the installation program.

When configuring the management network, select a DHCP-assigned address for now and record the IP address as described in ["Recording the Management IP Address" on page 50](#). (You can optionally specify a static IP address for each node later, after converting the second node.)



Caution: Do not convert the remaining node of the everRun MX system at this point; otherwise, all everRun MX data and VMs will be lost.

3. When you finish installing the everRun 7.x software on the first node, verify that you can connect to the everRun Availability Console at the IP address of the newly-installed node.
4. Log on to the everRun Availability Console of the newly-installed node as described in ["Logging on to the everRun Availability Console for the First Time" on page 54](#).

When prompted for the initial configuration settings, type the static IP address you obtained from your network administrator as the **System IP** address. Also, if you want to fully enable the features of your everRun 7.x system for testing, upload and activate your product license on the **LICENSE INFORMATION** page.

Notes:



- When specifying the **System IP** address, type the system-wide IP address of the everRun system, and not the node0 or node1 address.
- If you want to verify that your VMs work on the first node before you install the everRun 7.x software on the remaining node, activate your product license now. You can use the P2V client to migrate your VMs to the everRun 7.x system without a product license, but you cannot start and test your VMs on the everRun 7.x system unless you activate a valid license.

To migrate VMs from the everRun MX node to the everRun 7.x node

With the first node running the everRun 7.x software and the second node running the everRun MX software, do the following:

1. If applicable, prepare your VMs for migration as described in ["Migrating a Physical Machine or Virtual Machine to an everRun 7.x System" on page 151](#). (If you need to migrate a Windows Server 2003 VM, see the specific steps in ["Migrating a Windows Server 2003 VM to an everRun 7.x System" on page 179](#).)

In some cases, you need to perform steps in the guest operating system before migrating a VM to ensure that the VM will function properly on the everRun 7.x system.

2. Log on to the **everRun Availability Center** on the remaining node of the everRun MX system at:

`http://everRunMX-system:8080`

3. In the left-hand navigation panel, click **Virtual Machines**.
4. Right-click a VM that you want to migrate and click **Unprotect**.
5. When the VM is unprotected and automatically shut down, go back to **XenCenter**.
6. In the left-hand navigation panel of **XenCenter**, locate and expand the entry for the everRun MX system. Click the VM, and click **Start**.
7. After the VM starts, click the **Console** tab and click **Click here to create a DVD Drive**. Shut down the VM to save the change.
8. Insert the P2V client CD into the DVD drive of the remaining everRun MX node.

9. In the **Console** tab, next to **DVD drive *n***, select the physical P2V client CD in the drop down menu. Click **Start** to begin booting the VM from the P2V client CD.
10. Migrate the VM by following the steps in "[Migrating a Physical Machine or Virtual Machine to an everRun 7.x System](#)" on page 151.
11. After the migration completes, power off the VM and close the VM console window.
12. In the everRun Availability Console connected to the everRun 7.x node, verify that the VM appears on the **Virtual Machines** page.
13. Start the migrated VM and verify that it is functioning properly. Follow the instructions in "[Migrating a Physical Machine or Virtual Machine to an everRun 7.x System](#)" on page 151 to complete any migration steps in the VM. For example, you may need to install drivers or disable certain services.



Caution: The original VM on the everRun MX system must remain shut down when you use the VM on the everRun 7.x system; otherwise, the VMs will have network and software licensing conflicts.



Note: You can start a VM on the everRun 7.x system only if you have activated your product license. Upload and activate your license as described in "[Managing the everRun Product License](#)" on page 67.

14. If necessary, configure and manage your VM as described in "[Managing Virtual Machines](#)" on page 135. For guest-specific settings, see:
 - "[Configuring Windows-based Virtual Machines](#)" on page 215
 - "[Configuring Linux-based Virtual Machines](#)" on page 221
15. Follow steps 1-14 to migrate additional VMs.
16. Verify that all of your VMs are functioning properly, and that you have recorded any additional settings that you need from the remaining everRun MX server, which you will overwrite in the next procedure.

To complete the conversion to the everRun 7.x software



Caution: Converting a node to the everRun 7.x software erases all hard drives in that



node. After you convert the second node, you cannot recover the original VMs except by restoring from exports or third-party backups.

1. Shut down the remaining node of the everRun MX system.
2. Follow the instructions in to install the everRun 7.x software on the remaining node. Power on the node, update the necessary settings in the firmware (BIOS or UEFI) setup utility, and boot the node from the everRun 7.x DVD to run the installation program.

When configuring the management network, select a DHCP-assigned address for now. (You can specify a static IP address after the software installation.)

3. When the installation finishes, connect to the everRun Availability Console at the system IP address of the everRun 7.x system.
4. On the **Physical Machines** page, wait until both PMs reach the **running** state, and then assign logical disks to storage groups on the everRun 7.x system, as described in "[Assigning a Logical Disk to a Storage Group](#)" on page 124.

Notes:



- When the second PM joins the everRun system, the system automatically adds the secondary everRun system disk to the Initial Storage Group; however, the system does not assign any other logical disks from the second PM to existing storage groups.
- If you assigned logical disks to the Initial Storage Group or other storage groups on the first PM, you must manually add matching logical disks from the second PM to the same storage groups; otherwise, the everRun system cannot fully synchronize.

5. Verify that both PMs reach the **running** state and that the PMs finish synchronizing. The initial synchronization may take minutes or hours depending on your configuration, including the amount of storage and the number of VMs.
6. Optionally, update the network settings for the everRun 7.x system:
 - If you want to reuse the static IP address of the everRun MX system as the system IP address of the everRun 7.x system, open the **Preferences** page and click **IP**

Configuration. On the **System IP** tab, enter the static IP settings that were used by the everRun MX system and click **Save**.

- If you want to specify a static IP address for each node, click each **Node IP** tab, enter the new settings, and click **Save**.

If necessary, the everRun Availability Console reloads to reflect the new addresses.

7. Configure the everRun 7.x settings summarized in ["Post-Installation Tasks" on page 53](#).

Troubleshooting

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

To resolve network connectivity problems with the everRun 7.x system

If you have trouble connecting to the everRun Availability Console after installing the first node, you may have used the same IP address for node0 and for the system IP address of the everRun 7.x system. To correct the problem, reinstall the everRun 7.x software on node0 and ensure that the IP addresses you type for node0 and the system IP address are unique.

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 7.x system.

For all systems, see ["Creating and Migrating Virtual Machines" on page 141](#) for information about migrating your virtual machines (VMs) to the everRun 7.x system.



Note: For best results, contact your authorized Stratus service representative for assistance in evaluating and performing upgrades from an Avance unit.

Platform Requirements

Whether you reuse the existing Avance hardware or migrate to new hardware, the platform must meet the minimum system requirements for everRun systems, as described in ["Physical Machine System Requirements" on page 476](#).

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 software. See ["Compatible Guest Operating Systems" on page 474](#).

Also, verify that each guest operating system is supported by the migration process (as described in ["Migrating a Physical Machine or Virtual Machine to an everRun 7.x System" on page 151](#)) or the import process (as described in ["Importing an OVF File from an everRun MX System" on page 182](#)).

Network Preparation

Prepare the platform network and the networking environment to conform to the everRun system requirements. See ["General Network Requirements and Configurations" on page 27](#).

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 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 system.

Only one everRun 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 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 system. For physical storage requirements, see ["Storage Requirements" on page 26](#).

Installing everRun

After the nodes in the everRun system are configured, you can install and configure the everRun software, as described in ["Software Installation" on page 34](#).

Migrating Virtual Machines

Using either the P2V client migration process or the OVF import process, migrate the VMs to the everRun system. For an overview of each process, see ["Creating and Migrating Virtual Machines" on page 141](#).

Converting an Avance Unit to an everRun 7.x System

Convert an Avance unit to an everRun system to perform an in-place migration of the Avance unit and its virtual machines (VMs) to the everRun 7.x software.

To convert an Avance unit, shut down one physical machine (PM) or *node* in the Avance unit and install the everRun software on that node. Use the P2V client to transfer each VM from the Avance node to the everRun node over the network. Then, install the everRun software on the remaining node.



Caution: Consider backing up the Avance unit and its VMs and recording its settings before the conversion. Converting an Avance unit to an everRun system ultimately overwrites everything on your Avance unit (after you migrate your VMs to the everRun node).

Notes:



- For best results, contact your authorized Stratus service representative for assistance in evaluating and performing upgrades from an Avance unit.
- Before converting an Avance system to an everRun system, verify that your PMs and VMs are supported as described in "Physical Machine System Requirements" on page 476 and "Compatible Guest Operating Systems" on page 474.

To prepare for converting an Avance unit

1. Plan for converting your Avance unit by reviewing the following information:
 - "Planning to Migrate from an Avance Unit" on page 171
Describes some items to consider when migrating or converting from an Avance unit to an everRun system.
 - "Software Installation" on page 34
Summarizes the steps for installing the everRun software.
 - "Migrating a Physical Machine or Virtual Machine to an everRun 7.x System" on page 151
Describes how to use the P2V client to migrate a VM from one system to another. Also describes some steps you may need to complete in your guest operating systems **before** migrating the VMs to ensure that the VMs will function properly on the everRun system.
2. Back up your Avance unit and VMs.
3. Download the everRun ISO file from the **Downloads** page at <http://www.stratus.com/go/support/everrun>.
4. Download the P2V client ISO file from the **Drivers and Tools** section of the same support page.
5. Burn the everRun ISO file to a physical DVD that you will use to install the everRun software on each PM in your system.
6. In the Avance Management Console, use the P2V client ISO file to create a VCD that you will boot in each Avance VM to transfer the VMs to the everRun system.

7. Contact your network administrator to request at least one static IP address to use as the system-wide IP address for your converted everRun system. Request an additional static IP address for each of the two nodes if you do not have a DHCP server to automatically assign these addresses or if you prefer to use only static addresses.



Note: You must maintain unique system IP addresses for the Avance unit and the everRun system while both systems are online; however, if you want to reuse the original Avance unit IP address for the everRun system, you can change the network settings of the everRun system after the conversion is complete.

To convert node0 of the Avance unit to an everRun node



Caution: Converting a node to the everRun software erases all hard drives in that node.

Starting with both nodes running the Avance software, do the following:

1. In Avance Management Console, verify that your Avance unit is running properly and that both PMs are online.
2. Enable maintenance mode on **node0** of the Avance unit.



Note: Start with node0 of the Avance unit for consistency, because this first node that you convert will become node0 of the everRun system.

3. Verify that the VMs migrate from node0 to node1.
4. Shut down node0.
5. Follow the instructions in ["Installing Software on the First PM" on page 45](#) to install the everRun software on node0. Power on the node, update the necessary settings in the firmware (BIOS or UEFI) setup utility, and boot the node from the everRun DVD to run the installation program.

When configuring the management network, select a DHCP-assigned address for now and record the IP address as described in ["Recording the Management IP Address" on page 50](#). (You can optionally specify a static IP address for each node later, after converting the second node.)



Caution: Do not convert the remaining node of the Avance unit at this point; otherwise, all Avance data and VMs will be lost.

6. When you finish installing the everRun software on node0, verify that you can connect to the everRun Availability Console at the IP address of the newly-installed node.
7. Log on to the everRun Availability Console on node0 as described in "[Logging on to the everRun Availability Console for the First Time](#)" on page 54.

When prompted for the initial configuration settings, type the static IP address you obtained from your network administrator as the **System IP** address. Also, if you want to fully enable the features of your everRun system for testing, upload and activate your product license on the **LICENSE INFORMATION** page.

Notes:



- When specifying the **System IP** address, type the system-wide IP address, and not the node0 or node1 address.
- If you want to verify that your VMs work on node0 before you install the everRun software on the remaining node, activate your product license now. You can use the P2V client to migrate your VMs to the everRun system without a product license, but you cannot start and test your VMs on the everRun system unless you activate a valid license.

To migrate VMs from the Avance node to the everRun node

With node0 running the everRun software and node1 running the Avance software, do the following:

1. If applicable, prepare your VMs for migration as described in "[Migrating a Physical Machine or Virtual Machine to an everRun 7.x System](#)" on page 151. (If you need to migrate a Windows Server 2003 VM, see the specific steps in "[Migrating a Windows Server 2003 VM to an everRun 7.x System](#)" on page 179.)

In some cases, you need to perform steps in the guest operating system before migrating a VM to ensure that the VM will function properly on the everRun system.

2. In the Avance Management Console, shut down a VM that you want to migrate.

3. Boot the VM from the P2V client VCD and migrate the VM by following the steps in ["Migrating a Physical Machine or Virtual Machine to an everRun 7.x System" on page 151](#).
4. After the migration completes, power off the VM and close the VM console window.
5. In the everRun Availability Console connected to the everRun node, verify that the VM appears on the **Virtual Machines** page.
6. Start the migrated VM and verify that it is functioning properly. Follow the instructions in ["Migrating a Physical Machine or Virtual Machine to an everRun 7.x System" on page 151](#) to complete any migration steps in the VM. For example, you may need to install drivers or disable certain services.



Caution: The original VM on the Avance system must remain shut down when you use the VM on the everRun system; otherwise, the VMs will have network and software licensing conflicts.



Note: You can start a VM on the everRun system only if you have activated your product license. Upload and activate your license as described in ["Managing the everRun Product License" on page 67](#).

7. If necessary, configure and manage your VM as described in ["Managing Virtual Machines" on page 135](#). For guest-specific settings, see:
 - ["Configuring Windows-based Virtual Machines" on page 215](#)
 - ["Configuring Linux-based Virtual Machines" on page 221](#)
8. Follow steps 1-7 to migrate additional VMs.
9. Verify that all of your VMs are functioning properly, and that you have recorded any additional settings you that need from the remaining Avance node (node1), which you will overwrite in the next procedure.

To complete the conversion to the everRun software



Caution: Converting a node to the everRun software erases all hard drives in that node. After you convert the second node, you cannot recover the original VMs except by restoring from exports or third-party backups.

1. Shut down the Avance unit to power off the remaining Avance node (node1). In the Avance Management Console, on the **Unit** page, click **Shutdown** .
2. Follow the instructions in to install the everRun software on node1. Power on the node, update the necessary settings in the firmware (BIOS or UEFI) setup utility, and boot the node from the everRun DVD to run the installation program.

When configuring the management network, select a DHCP-assigned address for now. (You can specify a static IP address after the software installation.)

3. When the installation finishes, connect to the everRun Availability Console at the system IP address of the everRun system.
4. On the **Physical Machines** page, wait until both PMs reach the **running** state, and then assign logical disks to storage groups on the everRun 7.x system, as described in "[Assigning a Logical Disk to a Storage Group](#)" on page 124.

Notes:



- When the second PM joins the everRun system, the system automatically adds the secondary everRun system disk to the Initial Storage Group; however, the system does not assign any other logical disks from the second PM to existing storage groups.
- If you assigned logical disks to the Initial Storage Group or other storage groups on the first PM, you must manually add matching logical disks from the second PM to the same storage groups; otherwise, the everRun system cannot fully synchronize.

5. Verify that both PMs reach the **running** state and that the PMs finish synchronizing. The initial synchronization may take minutes or hours depending on your configuration, including the amount of storage and the number of VMs.
6. Optionally, update the network settings for the everRun system:
 - If you want to reuse the static IP address of the Avance unit as the system IP address of the everRun system, open the **Preferences** page and click **IP Configuration**. On the **System IP** tab, enter the static IP settings that were used by the Avance unit and click **Save**.

- If you want to specify a static IP address for each node, click each **Node IP** tab, enter the new settings, and click **Save**.

If necessary, the everRun Availability Console reloads to reflect the new addresses.

7. Configure the everRun settings summarized in "[Post-Installation Tasks](#)" on page 53.

Troubleshooting

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

To resolve network connectivity problems with the everRun system

If you have trouble connecting to the everRun Availability Console, especially after installing the first node (node0), you may have used the same IP address for node0 and for the system IP address. To correct the problem, reinstall the everRun software on node0 and ensure that the IP addresses you type for node0 and the system IP address are unique.

Migrating a Windows Server 2003 VM to an everRun 7.x System

Follow these instructions to migrate a Windows Server 2003 virtual machine (VM) from an Avance unit or an everRun MX system to a destination everRun 7.2 or later system. You should understand the following considerations before you migrate the Windows Server 2003 VM:

- The Windows Server 2003 operating system is no longer supported by Microsoft.
- The **only version** of Windows Server 2003 that everRun systems support is the **Windows Server 2003 R2 Enterprise SP2 32-bit** operating system.
- The destination system must be running everRun software release 7.2 or later.

To migrate the VM, boot the *P2V client (virt-p2v)* in the source Windows Server 2003 VM and use the client to configure, initiate, and monitor a secure network transfer from source side. To begin, follow the appropriate procedure: *To prepare for migrating a Windows Server 2003 VM* from your source system, and then continue with *To migrate a Windows Server 2003 VM from Avance or everRun MX system*.

To prepare for migrating a Windows Server 2003 VM from an Avance unit

1. Download the P2V client ISO file from the **Drivers and Tools** section of the **Downloads** page at <http://www.stratus.com/go/support/everrun>.
 - a. On the **Downloads** page, click **everRun** (if it is not already displayed) and then select the appropriate version.

- b. Scroll down to **Drivers and Tools** and then continue scrolling to **everRun P2V Client for Virtual or Physical Machine Migration**.
 - c. Select the **P2V Client (virt-p2v)** file.
2. If you want to verify the integrity of the ISO image, also download the associated `fciv` checksum file, and then download the Microsoft File Checksum Integrity Verifier (FCIV) executable file from the Microsoft Support Website. Save both files to the directory that contains the downloaded ISO file.

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

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

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

3. In the Avance Management Console, use the P2V client ISO file to create a VCD that you will boot in the Windows Server 2003 VM to transfer it to the everRun system.
4. On the **Virtual Machines** page, select the Windows Server 2003 VM and click **Shutdown**.
5. When the Windows Server 2003 VM is stopped, click **Boot from CD**.
6. In the **Boot from a CD** dialog box, select the P2V client VCD and click **Boot**.

To prepare for migrating a Windows Server 2003 VM from an everRun MX system

1. Download the P2V client ISO file from the **Drivers and Tools** section of the **Downloads** page at <http://www.stratus.com/go/support/everrun>.
2. Burn the P2V client ISO file to a physical CD that you will boot in the Windows Server 2003 VM to transfer it to the everRun 7.2 or later system.
3. Perform steps 1 through 9 in the *To migrate VMs from the everRun MX node to the everRun 7.x node* section of "[Converting an everRun MX System to an everRun 7.x System](#)" on page [164](#) to shut down the Windows Server 2003 VM and boot it from the P2V client CD.

To migrate a Windows Server 2003 VM from an Avance or everRun MX system

1. In the **virt-p2v** window, fill in the destination everRun system's host name (or host IP address) and password. Click **Test connection** and **Next**.

2. In the next **virt-p2v** window, click **Start conversion**.

You can monitor the progress of the migration in the **virt-p2v** window and on the **Volumes** page of the destination everRun system's everRun Availability Console as volumes associated with the new VM begin to appear.

3. When the migration is complete, the **virt-p2v** window displays a success message. Click **Power Off** to shut down the source VM.
4. In the destination everRun system's everRun Availability Console, click **Virtual Machines**.
5. Select the newly created VM and click **Start**.
6. Log on to the Windows Server 2003 guest operating system.
7. A service control manager warning about a driver failing during system startup appears. Click **OK**.
8. In the **Found New Hardware** wizard, select **No, not this time** and click **Next**.
9. Select **Install the software automatically**. Click **Next**.
10. A warning about the RedHat VirtIO Ethernet Adapter not passing Windows Logo testing appears. Click **Continue Anyway**.
11. When the **Found New Hardware** wizard completes, click **Finish**.
12. A warning about the RedHat VirtIO SCSI Adapter not passing Windows Logo testing appears. Click **Continue Anyway**.
13. The **Found New Hardware** wizard displays a **Cannot Install this Hardware** message. Select **Don't prompt me again to install this software** and click **Finish**.
14. At the prompt to restart the computer, click **Yes**.
15. A service control manager warning about a driver failing during system startup appears again. Click **OK**.
16. 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 system may continue to synchronize data between PMs to enable High Availability (HA) or Fault Tolerant (FT) operation.

Related Topics

["Migrating From Avance or everRun MX Systems" on page 160](#)

["Creating and Migrating Virtual Machines" on page 141](#)

["Configuring Windows-based Virtual Machines" on page 215](#)

["Managing Virtual Machine Resources" on page 230](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

Importing an OVF File from an everRun MX System

Import an Open Virtualization Format (OVF) file from an everRun MX system if you want to transfer a VM to the everRun 7.x system for deployment. (To migrate a physical machine (PM) or virtual machine (VM) to the everRun 7.x system without using an OVF file, see ["Migrating a Physical Machine or Virtual Machine to an everRun 7.x System" on page 151.](#))

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 7.x 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 Server 2008 from everRun MX systems. Importing a Windows Server 2003 VM from an OVF file is not supported. If you need to transfer a Windows Server 2003 VM to an everRun 7.x system, see ["Migrating a Windows Server 2003 VM to an everRun 7.x System"](#) on page 179.
- 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 7.x 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 the everRun Availability Console. You use XenConvert to export the VM to this share, then you import the VM to the everRun 7.x system from this share.
- In preparation for exporting the OVF file from the everRun MX system, you must unprotect the VM in the everRun Availability Center, which automatically shuts down the VM. 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 volume over a 1 Gb/s network may take about 30 minutes each way, export and import.
- When you import the VM on the everRun 7.x system, the import wizard creates a new instance of the VM with unique hardware IDs. The import wizard does not offer the Restore option that creates an identical VM with the same hardware IDs (SMBIOS UUID, system serial number, and MAC addresses), because the export files from everRun MX systems do not include this information.
- If you will continue to use the source VM on the everRun MX system after the import, remember to set a different IP address and hostname for the VM on the everRun 7.x system.
- If the everRun 7.x 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 7.x 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 Availability Center** with the hostname or IP address of your everRun MX master node at:

`http://everRunMX-system:8080`

2. In the left-hand navigation panel, click **Virtual Machines**.
3. Right-click a VM that you want to export and click **Unprotect**.
4. When the VM is unprotected and automatically shut down, open **Citrix XenCenter**.
5. In the left-hand navigation panel of **XenCenter**, locate and expand the entry for the everRun MX system. Click the VM that you want to export, and click **Start**.
6. Click the **Console** tab to open the console of the VM and log on to the Windows guest operating system.
7. Ensure that all volumes are labeled accurately, as summarized in "[Managing Windows Drive Labels](#)" on page 214.
8. Run the Windows System Preparation Tool (*Sysprep*) to prepare the guest operating system for redeployment.
9. 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 **Downloads** 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 7.x 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 7.x System

Importing a VM to the everRun 7.x 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 7.x system must be online for the import process to function properly.

To import a VM to the everRun 7.x 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 7.x system with the everRun Availability Console.
3. On the **Physical Machines** page (see "[The Physical Machines Page](#)" on page 95), 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 99), click **Import/Restore** to open the import wizard.
5. Click **Browse**. In the file browser, select the **.ovf** file that you want to import from your management PC and click **Import**.
6. Click **Import** to create a new instance of the VM with unique hardware IDs.

7. When prompted, click **Browse** and select a **.vhd** file to include for each volume associated with the VM.
8. 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 7.x 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.
9. Optionally, clear the check box for **Auto start Virtual Machine after import** if you need to reprovision the VM before starting it for the first time on the everRun 7.x system.
10. 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.

11. If applicable, use the **Reprovision Virtual Machine** wizard to allocate additional resources to the VM, as described in ["Reprovisioning Virtual Machine Resources" on page 231](#).
When you are finished reprovisioning the VM, click **Start** to boot the VM.
12. Click **Console** to open the console of the VM and log on to the guest operating system.
13. Update the VirtIO drivers to the latest supported versions, as described in ["Updating the VirtIO Drivers \(Windows-based VMs\)" on page 215](#).
14. If necessary, update the network settings in the guest operating system.

After you verify that the new VM is functioning properly, the import process is complete; however, the everRun 7.x 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 7.x 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 7.x 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 7.x 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 7.x system

If your data volumes do not appear in the VM on the everRun 7.x 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 7.x 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 Avance or everRun MX Systems" on page 160](#)

["Creating and Migrating Virtual Machines" on page 141](#)

["Configuring Windows-based Virtual Machines" on page 215](#)

["Configuring Linux-based Virtual Machines" on page 221](#)

["Managing Virtual Machine Resources" on page 230](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

Importing an OVF File from an Avance System

Import an Open Virtualization Format (OVF) file from an Avance unit if you want to transfer a VM to the everRun 7.x system for deployment. (To migrate a physical machine (PM) or virtual machine (VM) to the everRun 7.x system without using an OVF file, see ["Migrating a Physical Machine or Virtual Machine to an everRun 7.x System" on page 151.](#))

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 system.

When you import a VM image in the everRun Availability Console, the import wizard allows you to choose between *importing* or *restoring* the VM. Importing a VM creates a new instance of the VM with unique hardware IDs. Restoring a VM creates an identical VM with the same hardware IDs (SMBIOS UUID, system serial number, and MAC addresses, if provided in the VM image) that your guest operating system and applications may require for software licensing. To prevent conflicts with the original VM, restore a VM only if you want to transfer it to the everRun system and stop using it on the source 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.
- If you need to transfer a Windows Server 2003 VM to an everRun system, see "[Migrating a Windows Server 2003 VM to an everRun 7.x System](#)" on page 179. Importing a Windows Server 2003 VM from an OVF file is not supported.
- 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 7.x 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 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 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 volume over a 1 Gb/s network may take about 30 minutes each way, export and import.





- To prevent conflicts with the source VM on the Avance unit, the import wizard automatically assigns a new MAC address to each network interface in the new VM; however, you may need to manually update any IP addresses and host names as needed.
- If the everRun 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 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 214.
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 **Downloads** 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 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 System

Importing a VM to the everRun 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 system must be online for the import process to function properly.

To import a VM to the everRun system

1. Log on to the everRun system with the everRun Availability Console.
2. On the **Physical Machines** page (see "[The Physical Machines Page](#)" on page 95), 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 99), click **Import/Restore** to open the import wizard.
4. Click **Browse**. In the file browser, select the **.ovf** file that you want to import from your management PC and click **Import**.
5. Select **Import** or **Restore**. Import creates a new instance of the VM with unique hardware IDs. Restore creates an identical VM with the same hardware IDs that are provided in the OVF file.
6. When prompted, click **Browse** and select a **.vhd** file to include for each volume associated with the VM.
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 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.

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 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 231.
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 only, update the VirtIO drivers to the latest supported versions, as described in "[Updating the VirtIO Drivers \(Windows-based VMs\)](#)" on page 215.
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 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 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 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 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 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 227](#).

To recover missing data volumes in the VM on the everRun system

If your data volumes do not appear in the VM on the everRun 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 (`/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 system

If your network devices do not appear in the VM on the everRun 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 Avance or everRun MX Systems" on page 160](#)

["Creating and Migrating Virtual Machines" on page 141](#)

["Configuring Windows-based Virtual Machines" on page 215](#)

["Configuring Linux-based Virtual Machines" on page 221](#)

["Managing Virtual Machine Resources" on page 230](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

Importing an OVF File from an everRun 7.x System

Import an Open Virtualization Format (OVF) file from an everRun system if you want to transfer a VM from one everRun 7.x system to another, or if you want to transfer an image that you created back to the same everRun 7.x system to restore or duplicate the original VM. (To migrate a physical machine (PM) or virtual machine (VM) to the everRun 7.x system without using an OVF file, see ["Migrating a Physical Machine or Virtual Machine to an everRun 7.x System" on page 151.](#))

To import a VM image from an everRun system, first use the everRun Availability Console on the source everRun system to export a VM (["Exporting a Virtual Machine from an everRun System" on page 208](#)) or a VM snapshot (["Exporting a Snapshot" on page 256](#)) to OVF and Virtual Hard Disk (VHD) files on a supported network share. Copy these files to your management PC (or mount the network share), and then

use the everRun Availability Console on the target everRun system to import the OVF and VHD files from your management PC.

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

If you want to restore an existing VM on the same everRun system to overwrite the VM and recover it from a previous backup copy, see ["Replacing a Virtual Machine from an OVF File" on page 203](#).

Notes:

- You can import VMs only if they are running supported guest operating systems, as described in ["Compatible Guest Operating Systems" on page 474](#).
- The time required for the 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 volume over a 1 Gb/s network may take about 30 minutes.
- When you import or restore an everRun VM, the original container size for each volume you include is not preserved. For example, if your source VM has a 20 GB boot volume in a 40 GB volume container, the target VM will have a 20 GB boot volume in a 20 GB volume container. If necessary, you can expand the volume containers on the target system as described in ["Expanding a Volume Container on the everRun System" on page 241](#).
- If you are importing a VM back to the same everRun system to duplicate the VM, you must rename the VM and duplicate volumes during either the export or import process.
- To prevent conflicts with the source VM, the import wizard automatically assigns a new MAC address to each network interface in the new VM; however, you may need to manually update any IP addresses and host names as needed.
- If the everRun 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 system, and import them again.



Prerequisite: Both PMs of the everRun system must be online for the import process to function properly.

To import a VM to the everRun system

1. Export a VM or a VM snapshot from the source everRun system.
2. Copy the exported OVF and VHD files to the management PC that you use to run the everRun Availability Console, or mount the export network share that contains these files.
3. Log on to the everRun Availability Console on the target everRun system.

4. On the **Physical Machines** page (see ["The Physical Machines Page" on page 95](#)), verify that both PMs are in the **running** state and that neither PM is in maintenance mode or in the process of synchronizing.
5. On the **Virtual Machines** page (see ["The Virtual Machines Page" on page 99](#)), click **Import/Restore** to open the import wizard.
6. Click **Browse**. In the file browser, locate the exported files. Select the **.ovf** file that you want to import, and click **Import**.
7. Select **Import** or **Restore**. Import creates a new instance of the VM with unique hardware IDs. Restore creates an identical VM with the same hardware IDs that are provided in the OVF file.
8. When prompted, click **Browse** and select a **.vhd** file to include for each volume associated with the VM.
9. 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 system (the boot volume is required). Select the **Restore Data** box to import data for a volume from the VHD file.

If the target everRun system has more than one storage group, you can also select the storage group in which to create each volume. Ensure that you select a **Destination** group that supports the sector size of the volume you are importing (see ["Planning Virtual Machine Storage" on page 139](#)) and select the **Sector Size** that matches the source volume (the import wizard cannot convert the sector size of a volume). Note that the boot volume must have a sector size of 512 B. You can select the sector size, either 4K or 512B, only for data disks.

- **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.

10. 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 system.
11. 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.

12. If applicable, use the **Reprovision Virtual Machine** wizard to allocate additional resources to the VM, as described in "[Reprovisioning Virtual Machine Resources](#)" on page 231. Also, if you want to allocate additional space in each volume container for snapshots, see "[Expanding a Volume Container on the everRun System](#)" on page 241.

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

13. Click **Console** to open the console of the VM and log on to the guest operating system.
14. For Windows-based VMs only, download and update the VirtIO drivers to the latest supported versions, as described in "[Updating the VirtIO Drivers \(Windows-based VMs\)](#)" on page 215. (The correct VirtIO drivers are already present in Linux-based VMs.)



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

15. 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 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.

Troubleshooting

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

To clean up after a canceled or failed import

In the everRun Availability Console on the target system, remove the imported VM and any volumes associated with the imported VM.

To recover missing data volumes in the target VM

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

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

To recover missing network devices in the VM on the everRun system

If your network devices do not appear in the VM on the target everRun 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 Avance or everRun MX Systems" on page 160](#)

["Creating and Migrating Virtual Machines" on page 141](#)

["Configuring Windows-based Virtual Machines" on page 215](#)

["Configuring Linux-based Virtual Machines" on page 221](#)

["Managing Virtual Machine Resources" on page 230](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

Replacing a Virtual Machine from an OVF File

Replace a virtual machine (VM) from an Open Virtualization Format (OVF) file if you want to restore a VM on your everRun system by overwriting the VM with a previous backup copy. (If you want to import a VM from a different system, see the overview in ["Creating and Migrating Virtual Machines" on page 141.](#))

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

You can restore a VM that already exists on an everRun system only if you have previously exported a VM (["Exporting a Virtual Machine from an everRun System" on page 208](#)) or a VM snapshot (["Exporting a Snapshot" on page 256](#)) to OVF and Virtual Hard Disk (VHD) files on a supported network share. Copy these files to your management PC (or mount the network share), and then use the everRun Availability Console on the target everRun system to restore the OVF and VHD files from your management PC.



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

Notes:

- The time required to restore a VM depends on the size and number of volumes in the source VM as well as network bandwidth. For example, transferring a VM with one 20 GB boot volume over a 1 Gb/s network may take about 30 minutes.
- If you overwrite and restore an existing VM, the everRun system removes the existing VM and its volumes, but the system does not remove any of the VM's snapshots or the volume containers in which the snapshots are stored. The volume containers continue to use storage space on your everRun system until you remove the VM's snapshots (see ["Removing a Snapshot" on page 262](#)). If storage space is limited, you may want to remove the snapshots before starting the restore process to ensure that there will be enough storage space for the operation.
- If you previously expanded the volume containers of your VM to allow enough space for snapshots, you may want to note the current size of each volume container before you overwrite and restore the VM. Because the everRun system creates all new volume containers for a restored VM and does not preserve the expanded container sizes, you need to manually expand the volume containers of the restored VM after the restore process is finished (see ["Expanding a Volume Container on the everRun System" on page 241](#)).
- If the everRun system switches from the primary PM to the secondary PM while restoring a VM, the restore process fails. This does not affect the continuous uptime of your system, but you must delete the incomplete VM and associated volumes on the everRun system, and restore them again.



Prerequisite: Both PMs of the everRun system must be online for the restore process to function properly.

To overwrite and restore a VM on the everRun system

1. Export a VM or a VM snapshot from the source everRun system.
2. Copy the exported OVF and VHD files to the management PC that you use to run the everRun Availability Console, or mount the export network share that contains these files.
3. Log on to the everRun Availability Console on the target everRun system.

4. On the **Physical Machines** page (see ["The Physical Machines Page" on page 95](#)), verify that both PMs are in the **running** state and that neither PM is in maintenance mode or in the process of synchronizing.
5. On the **Virtual Machines** page (see ["The Virtual Machines Page" on page 99](#)), select the VM that you want to restore from your previous backup copy.
6. Click **Restore** to open the restore wizard.
7. Click **Browse**. In the file browser, locate the network share with the exported files. Select the **.ovf** file that you want to restore, and click **Restore**.
8. When prompted, click **Browse** and select a **.vhd** file to include for each volume associated with the VM.
9. Confirm that you will overwrite the existing VM and data and proceed by clicking **Continue**.



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

10. 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 system (the boot volume is required). Select the **Restore Data** box to import data for a volume from the VHD file.

If the target everRun system has more than one storage group, you can also select the storage group in which to create each volume. Ensure that you select a **Destination** group that supports the sector size of the volume you are importing (see ["Planning Virtual Machine Storage" on page 139](#)) and select the **Sector Size** that matches the source volume (the restore wizard cannot convert the sector size of a volume). Note that the boot volume must have a sector size of 512 B. You can select the sector size, either 4K or 512B, only for data disks.
 - **Network**

Displays all of the available networks. You can remove a network or add one that is not already allocated. A minimum of one network is required.

11. Optionally, clear the check box for **Auto start Virtual Machine after restore** if you need to re-provision the VM before starting it for the first time.
12. Click **Restore** to begin restoring the VM. When the transfer is complete, click **Done** to close the restore wizard.



Note: Restored volumes begin to appear on the **Volumes** page of the everRun Availability Console while the restore process is still in progress. Do not attach or remove any of these restored volumes until the restore window reports that the process is complete; otherwise, the restore process fails.

13. If applicable, use the **Reprovision Virtual Machine** wizard to allocate additional resources to the VM, as described in ["Reprovisioning Virtual Machine Resources" on page 231](#). Also, if you want to allocate additional space in each volume container for snapshots, see ["Expanding a Volume Container on the everRun System" on page 241](#).

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

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



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

Troubleshooting

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

To clean up after a canceled or failed restore process

In the everRun Availability Console on the target system, remove the restored VM and any volumes associated with the restored VM.

Related Topics

["Creating and Migrating Virtual Machines" on page 141](#)

["Managing Virtual Machine Resources" on page 230](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

Exporting a Virtual Machine from an everRun System

Export a virtual machine (VM) to save an image of the VM to a network-mounted folder (that is, directory). Exporting a VM from an everRun system makes the VM image available for importing to another system or for importing back to the same everRun system to restore or duplicate the original VM. (You can directly export a VM from the everRun system as described in this topic, or create and export a snapshot. For an overview of snapshots, see ["Managing Snapshots" on page 249.](#))

You prepare for exporting a VM by creating a network-mounted folder to store an exported VM in your environment. Create a folder for either a Windows share (also known as a Common Internet File System (CIFS) share) or a Network File System (NFS) export. After creating the folder, you mount it in the host operating system of the everRun system, as described in this topic. When you initiate an export in the everRun Availability Console, the everRun system saves the VM image to the network-mounted folder as standard Open Virtualization Format (OVF) and Virtual Hard Disk (VHD) files.

Notes:

- For Linux-based VMs, before exporting the VM, 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 may use different device names on another system, the 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 running some older Ubuntu releases, you may need to edit the `/boot/grub/grub.cfg` file and change the `gfxmode` parameter to `text` (for example, set `gfxmode=text`) before exporting a VM; otherwise, the new VM's console may hang on another system. You can restore the original setting in the source VM after the migration.
- The source VM must be shut down during the export process. Consider scheduling a planned maintenance period for this process (or consider taking a snapshot that you can export at a later time, as described in ["Creating a Snapshot" on page 250](#)).
- The time required for the export depends on the size and number of volumes in the source VM as well as network bandwidth. For example, transferring a VM with one 20 GB boot disk over a 1 Gb/s network may take about 30 minutes.
- If you will continue to use the source VM after exporting it, remember to set a different MAC address and IP address for the VM when you import it on the target system.
- If the everRun system switches from the primary PM to the secondary PM during an export, the process fails. This does not affect the continuous uptime of the system. You can delete the partially exported files from the network-mounted folder and export the files again.

**Prerequisites:**

- You must shut down a VM before exporting it.
- Both PMs of the everRun system must be online for the export process to function properly.

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

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

1. Log on to the everRun system with the everRun Availability 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 214.
5. Run the Windows System Preparation Tool (`Sysprep`) to prepare the guest operating system for redeployment.

To create and mount an NFS export or Windows CIFS share

Before you can export a VM, you must create and mount the folder to which the export(s) will be transferred, as follows:

1. Create a network-mounted folder for a Windows/CIFS share or an NFS export in your environment where you can store the exported VM.

Ensure that you have enough storage for the VMs that you want to export. Also, set full read/write permissions on the network-mounted folder to permit file transfers, or, for a Windows/CIFS share only, assign read/write permissions to a specific user on the system/domain that hosts the share. Record the location and settings of the NFS export or CIFS share, which you will specify in a later step.

2. Log on to the everRun system with the everRun Availability Console.
3. On the **Physical Machines** page, verify that both PMs are in the running state and that neither PM is in maintenance mode or in the process of synchronizing. See "[The Physical Machines Page](#)" on page 95. Also, make note of which PM is the **primary** node. In the list in the top pane, the primary node is marked as **noden (primary)**.
4. Obtain the IP address of the primary node, if you do not already know it. For example, on the **Preferences** page, click **IP Configuration**. Click the **Noden IP** tab for the primary node and record the **IP Address** value.
5. Use a secure shell (`ssh`) utility to log on to the host operating system (host OS) of the primary node of the everRun system, where you will mount the network share. Log on as the

root user.

The next step shows how to use the `ftxmnt` script to automatically mount the export share. The script should work in most cases, but you can also manually mount the share by executing standard mount commands if needed.

Notes:

- If you manually mount the NFS export or Windows/CIFS share in the host OS of the everRun system, you must create the mount point at `/mnt/ft-export`, where the export process expects to find it. (If you use the `ftxmnt` script, it automatically creates this mount point.)
- If you want the mount to persist when you reboot the everRun system, manually add an entry to the `/etc/fstab` file in the host OS of the everRun system. (The `ftxmnt` script does not modify the `/etc/fstab` file.)
- You need to mount the NFS export or Windows/CIFS share only on the primary node of the everRun system, but you can optionally add the mount to the `/etc/fstab` file on both nodes so it is always available even if the primary node changes.
- If you mount a Windows/CIFS share that requires a password, and the password changes after the mount, you need to unmount the share and mount it again with the new password; otherwise, the export may fail unexpectedly.



6. To automatically mount the NFS export or Windows/CIFS share, execute the `ftxmnt` script and respond to the interactive prompts. For example, the following output shows how to mount a Windows/CIFS share (`\\192.168.0.111\ExportVMs`) that is accessible by a specific user account:

```
[root@node0 /]# ftxmnt
```

```
This script is meant to mount a Network Attached Storage
location to use for exporting everRun virtual machines.
Enter Ctrl-C to exit
```

Enter **n** if you are mounting an nfs share, enter **w** if you are entering a windows share: **w**

What is the IP address or the computer name of the file server?

192.168.0.111

What is the name of the share you wish to mount?

ExportVMs

Does this share require authentication? (y/n):

Respond **y** to the prompt above if the share requires authentication. When you respond **y**, the following prompts appear:

- What is your username?

Enter your *domain\username*.

- Password for *domain\username*:

Enter your password.

If the share does not require authentication, type **n** and then respond to the following prompt:

If there is a guest account name, enter it here:

Type the guest account name, if it exists.

The following message appears:

```
Successfully mounted folder \\192.168.0.111\ExportVMs at
path /mnt/ft-export/
```

7. Switch to the `/mnt/ft-export` directory in the everRun host OS and create a file to verify that the export or share is present and that you have read/write permissions. For example:

```
# touch test
```

```
# ls
```

```
test
```

Confirm that the file also appears in the export or share on the remote system. If the file is not present or the everRun host OS displays an error, verify your mount settings and permissions.

8. Remove the test file.

```
# rm test
```

Later, you want to unmount the export or share after exporting your virtual machines, switch out of the `/mnt/ft-export` directory and execute the `umount` command, as follows:

```
# cd /
# umount /mnt/ft-export
```

To export a VM from the everRun system

1. Log on to the everRun system with the everRun Availability Console.
2. On the **Physical Machines** page (see ["The Physical Machines Page" on page 95](#)), 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 99](#)), select the VM you want to export and click **Shutdown**. Wait for the VM to shut down.
4. With the VM selected, click **Export** to open the export wizard.
5. In the **Export VM** dialog box, type the path in `/mnt/ft-export` where you want the VM to be exported and its OVF and VHD files to be stored. For example, if you want to export the VM to a new folder named `ocean1`, type `ocean1`. If you want the new `ocean1` folder to be inside an existing directory named `TestVMs`, type `TestVMs/ocean1`.
6. Select the volumes to include under **Data Volumes to Export**. (The boot volume is required.)
7. Click **Export VM**.

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

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

After the export process, if you want to import or restore the OVF and VHD files on an everRun system, see ["Importing an OVF File from an everRun 7.x System " on page 198](#)

Troubleshooting

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

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

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

Related Topics

["Creating and Migrating Virtual Machines" on page 141](#)

["Managing Virtual Machine Resources" on page 230](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

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 or create a snapshot of it.



Caution: Ensure that each volume has a unique identifiable label before running **Sysprep** (to prepare for an export or snapshot). 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 or snapshot will help to identify the drives. For instructions on reassigning drive letters on a Windows system, search for the Microsoft Support web site.

Related Topics

["Creating and Migrating Virtual Machines" on page 141](#)

["Configuring Windows-based Virtual Machines" on page 215](#)

["Managing Snapshots" on page 249](#)

Configuring Windows-based Virtual Machines

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

- ["Updating the VirtIO Drivers \(Windows-based VMs\)" on page 215](#)
- ["Creating and Initializing a Disk \(Windows-based VMs\)" on page 218](#)
- ["Installing Applications \(Windows-based VMs\)" on page 219](#)

If you plan to create VM snapshots (see ["Managing Snapshots" on page 249](#)), consider installing the QEMU guest agent and configuring the Microsoft Shadow Volume Copy Service (VSS), as described in:

- ["Installing the QEMU Guest Agent for Application-Consistent Snapshots \(Windows-based VMs\)" on page 219](#)

In addition, ensure that you configure the following settings:

- Change the time zone in the guest operating system to correspond to the time zone configured on the **Date and Time** preference page in the everRun Availability Console (see ["Configuring Date and Time" on page 75](#)); 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 system.
- Disable hibernation (enabled by default in some cases) to prevent the guest operating system from going into a power-saving state.
- Configure the power button action in the guest operating system to shut down the guest (and not to hibernate it) to allow the **Shutdown** VM button in the everRun Availability Console to work properly (see ["Shutting Down a Virtual Machine" on page 225](#)).
- Configure the guest operating system to generate a crash dump file if the operating system crashes. Follow the instructions in the Microsoft article, [How to generate a complete crash dump file or a kernel crash dump file by using an NMI on a Windows-based system](#) (Article ID: 927069). Follow the instructions in the **More Information** section.

Related Topics

["Managing Virtual Machines" on page 135](#)

Updating the VirtIO Drivers (Windows-based VMs)

Update the Red Hat VirtIO drivers in your Windows-based virtual machines (VMs) to the latest supported versions to ensure the proper operation of your VMs. For example, you should update the VirtIO drivers after upgrading the everRun software ("[Upgrading everRun Software](#)" on page 115) or after using the P2V client to migrate a VM to the everRun system ("[Migrating a Physical Machine or Virtual Machine to an everRun 7.x System](#)" on page 151).

Notes:



- For proper operation, ensure that you download the VirtIO drivers from the **everRun Support** page, as described in the following procedure. The support page contains a VirtIO ISO file with versions of the VirtIO drivers that have been tested with the everRun software.
- After updating the drivers, you may need to restart the guest operating system.

To update the VirtIO drivers in a Windows-based virtual machine

1. Download the VirtIO ISO file from the **Downloads** page at <http://www.stratus.com/go/support/everrun>.
 - a. On the **Downloads** page, click **everRun** (if it is not already displayed) and then select the appropriate version.
 - b. Scroll down to **Drivers and Tools** and then continue scrolling to **everRun VirtIO Driver Update**.
 - c. Click the link to the appropriate file.

Ensure that you download the version of the VirtIO ISO file that matches the version of your everRun system.

2. If you want to verify the integrity of the ISO image, also download the associated `fciv` checksum file, and then download the Microsoft File Checksum Integrity Verifier (FCIV) executable file from the Microsoft Support web site. Save both files to the directory that contains the downloaded ISO file.

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

```
fciv -v -xml virtio-win-n.n.nn.xml
```

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

3. Open the everRun Availability Console, create a VCD of the VirtIO ISO file, and insert the VCD into the Windows-based VM (see ["Creating a Virtual CD" on page 244](#) and ["Inserting a Virtual CD" on page 246](#)).

4. Open **Device Manager** in the guest operating system.

5. Expand **Network adapters** and locate the **Red Hat VirtIO Ethernet Adapter**. There may be more than one adapter present depending on the number of network interfaces in your VM.

If the **Red Hat VirtIO Ethernet Adapter** is not present, the VirtIO driver is not installed. Expand **Other devices** and locate the unknown **Ethernet Controller** device. You will update the driver for this device.

6. Right click the **Red Hat VirtIO Ethernet Adapter** (or **Ethernet Controller**) and select **Update Driver Software**. Click **Browse my computer for the driver software**, specify the location of the VirtIO Ethernet driver (**netkvm**) for your guest operating system, and finish updating the driver. (For example, to update the driver in a Windows Server 2012 R2 guest, select the `NetKVM\2k12R2\amd64\netkvm.inf` file on the VirtIO VCD.)

Repeat the driver update for each additional **Red Hat VirtIO Ethernet Adapter** (or **Ethernet Controller**) device.

7. Expand **Storage controllers** and locate the **Red Hat VirtIO SCSI controller**. There may be more than one controller present depending on the number of volumes in your VM.

If the **Red Hat VirtIO SCSI controller** is not present, the VirtIO driver is not installed. Locate the unknown **SCSI controller** device. You will update the driver for this device.

8. Right-click the **Red Hat VirtIO SCSI controller** (or **SCSI controller**) and select **Update Driver Software**. Click **Browse my computer for the driver software**, specify the location of the VirtIO SCSI driver (**viostor**) for your guest operating system, and finish updating the driver. (For example, to update the driver in a Windows Server 2012 R2 guest, specify the `viostor\2k12R2\amd64\viostor.inf` file on the VirtIO VCD.)

Repeat the driver update for each additional **Red Hat VirtIO SCSI** (or **SCSI controller**) device.



Caution: Although the device name is the **Red Hat VirtIO SCSI** controller, you must select the storage driver file that is labeled **viostor**, and not **vioscsi** (if present).

Installing the **vioscsi** driver may crash your VM.

9. If you intend to use the QEMU guest agent, as described in ["Installing the QEMU Guest Agent for Application-Consistent Snapshots \(Windows-based VMs\)"](#) on page 219, also update the VirtIO serial driver that is associated with the guest agent; otherwise, skip to step 11.

Expand **System devices** and locate the **VirtIO Serial Driver**. If the **VirtIO Serial Driver** is not present, expand **Other devices** and locate the unknown **PCI Simple Communications Controller** device. You will update the driver for this device.

10. Right-click the **VirtIO Serial Driver** and select **Update Driver Software**. Click **Browse my computer for the driver software**, specify the location of the VirtIO serial driver (**vioser**) for your guest operating system, and finish updating the driver. (For example, to update the driver in a Windows Server 2012 R2 guest, specify the `vioserial\2k12R2\amd64\vioser.inf` file on the VirtIO VCD.)
11. If applicable, restart the guest operating system to load the updated drivers.

Related Topics

["Configuring Windows-based Virtual Machines" on page 215](#)

["Creating and Migrating Virtual Machines" on page 141](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

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 system, as described in ["Creating a Volume in a Virtual Machine" on page 234](#).
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 227](#)

["Configuring Windows-based Virtual Machines" on page 215](#)

["Creating and Migrating Virtual Machines" on page 141](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

Installing Applications (Windows-based VMs)

Install an application in a Windows-based virtual machine by doing one of the following:

- Download the installation program to the guest operating system as an executable file or ISO file.
- Mount a network drive that contains the installation program.
- Create and insert a Virtual CD (VCD) that contains the installation program. See ["Managing Virtual CDs" on page 244](#).

Related Topics

["Opening a Virtual Machine Console Session" on page 227](#)

["Configuring Windows-based Virtual Machines" on page 215](#)

["Creating and Migrating Virtual Machines" on page 141](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

Installing the QEMU Guest Agent for Application-Consistent Snapshots (Windows-based VMs)

Install the Quick EMUlator (QEMU) guest agent in your Windows-based guest operating system if you want to create application-consistent snapshots of your virtual machine (VM). For an overview of everRun snapshots, see ["Managing Snapshots" on page 249](#).

Typically, while applications are running, they process transactions, open and write files, hold information in memory, and more. If you take a VM snapshot while your applications are still working, it is similar to restarting your system after a power outage. Although most modern file systems are designed to recover from this type of outage, it is possible that some data will be corrupted or lost in the process, especially while transaction-intensive applications are running. In this case, taking a snapshot without preparing your applications results in a *crash-consistent* snapshot, as if you took the snapshot after a crash or power outage.

Microsoft Windows provides the Volume Shadow Copy Service (VSS) that can inform the file system and your applications when they must temporarily *quiesce* or freeze their operations during a snapshot or backup. If your applications support VSS, the everRun software can signal your applications through the QEMU guest agent and VSS to quiesce during a snapshot on your everRun system, thus ensuring an application-consistent snapshot.



Caution: Before installing the QEMU guest agent, contact your application vendor(s) to determine if they support Microsoft VSS and if any additional configuration steps are necessary to support VSS operations. You can create application-consistent snapshots only if your applications support VSS and the QEMU guest agent is properly installed and running.



Notes:

- By default, all snapshots are considered crash-consistent snapshots unless you install the QEMU guest agent and explicitly configure your applications to quiesce when signaled by Microsoft VSS.
- When you install the QEMU guest agent, you may need to restart your VMs. If your VMs are in use, schedule a maintenance period for this procedure.

To install the QEMU guest agent

1. Log on to the everRun system with the everRun Availability Console.
2. Select a VM on the **Virtual Machines** page.
3. Click **Console** and log on to the Windows guest operating system.

4. To transfer the QEMU guest agent installer to your system, do one of the following:
 - Open a web browser and download the installer from the **Drivers and Tools** section of the **Downloads** page at <http://www.stratus.com/go/support/everrun>.
 - Mount a local network share that contains the installer and either copy it to your system or prepare to run it from the share.
5. Start the installer by double-clicking the icon. The QEMU Guest Agent Setup wizard is displayed.
6. Read the license information. If appropriate, click the check box next to **I agree to the license terms and conditions**.
7. Click **Install** to begin the software installation.
8. If Windows prompts that it cannot verify the publisher of the driver software, click **Install** to continue installing the software.
9. If prompted, click **Restart** to restart the guest operating system.

When Windows restarts, you may see a message indicating that the driver software was installed.
10. If prompted, click **Restart** to restart the guest operating system again.

To verify that the QEMU guest agent is properly installed and running

Open **Services**. For example, click **Start** and **Run**, then type **services.msc** and click **Run**. Verify that the following services are present and running:

- QEMU Guest Agent (always runs)
- QEMU Guest Agent VSS Provider (may run only during quiesce)

Open **Device Manager**. For example, click **Start**, **Control Panel**, **Hardware**, and **Device Manager**.

Verify that the following driver is installed and running:

- VirtIO Serial Driver (under System devices)

Related Topics

["Configuring Windows-based Virtual Machines" on page 215](#)

["Creating and Migrating Virtual Machines" on page 141](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

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 222](#)
- ["Installing Applications \(Linux-based VMs\)" on page 223](#)

If you plan to create VM snapshots (see ["Managing Snapshots" on page 249](#)), consider installing the QEMU guest agent, as described in:

- ["Installing the QEMU Guest Agent for Application-Consistent Snapshots \(Linux-based VMs\)" on page 223](#)

In addition, ensure that you configure the following settings:

- Disable hibernation (enabled by default in some cases) to prevent the guest operating system from going into a power-saving state.
- Configure the power button action in the guest operating system to shut down the guest (and not to hibernate it) to allow the **Shutdown** VM button in the everRun Availability Console to work properly. For the minimal server version of **Ubuntu Linux**, optionally install the `acpid` package to enable the **Shutdown** button. See ["Shutting Down a Virtual Machine" on page 225](#).
- Install the `kexec-tools` package and configure the guest operating system to generate a crash dump file if the system crashes.
- For **Ubuntu Linux** guest operating systems, to prevent a problem where the VM console hangs in 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 before you can set the parameter, see the troubleshooting information in ["Opening a Virtual Machine Console Session" on page 227](#) to resolve the issue.

For more information about these settings, see your Linux documentation.

Related Topics

["Managing Virtual Machines" on page 135](#)

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 234](#).
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 227](#)

["Configuring Linux-based Virtual Machines" on page 221](#)

["Creating and Migrating Virtual Machines" on page 141](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

Installing Applications (Linux-based VMs)

Install an application in a Linux-based virtual machine by doing one of the following:

- Download the installation package to the guest operating system as an executable file or ISO file.
- Mount a network drive that contains the installation package.
- Create and insert a Virtual CD (VCD) that contains the installation package. See ["Managing Virtual CDs" on page 244](#).

Related Topics

["Opening a Virtual Machine Console Session" on page 227](#)

["Configuring Linux-based Virtual Machines" on page 221](#)

["Creating and Migrating Virtual Machines" on page 141](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

Installing the QEMU Guest Agent for Application-Consistent Snapshots (Linux-based VMs)

Install the Quick EMUlator (QEMU) guest agent in your Linux-based guest operating system if you want to create application-consistent snapshots of your virtual machine (VM). For an overview of everRun snapshots, see ["Managing Snapshots" on page 249](#).

Typically, while applications are running, they process transactions, open and write files, hold information in memory, and more. If you take a VM snapshot while your applications are still working, it is similar to restarting your system after a power outage. Although most modern file systems are designed to recover from this type of outage, it is possible that some data will be corrupted or lost in the process, especially for transaction-intensive applications. In this case, taking a snapshot without preparing your applications results in a *crash-consistent* snapshot, as if you took the snapshot after a power outage.

If your applications support QEMU signaling, the everRun software can signal your applications through the QEMU guest agent to ensure that your applications *quiesce* or freeze during a snapshot on your everRun system, thus ensuring an application-consistent snapshot.

Most Linux distributions already include a QEMU guest agent (usually in the `qemu-guest-agent` package). For information about installing and configuring the QEMU guest agent, see the documentation for your particular Linux distribution.



Caution: Before installing a QEMU guest agent, contact your application vendor(s) to determine if they support QEMU signaling and if any additional configuration steps are necessary to quiesce your applications. You can create application-consistent snapshots only if your applications support QEMU signaling and the QEMU guest agent is properly installed and running.



Notes:

- By default, all snapshots are considered crash-consistent snapshots unless you explicitly install the QEMU guest agent and configure your applications to quiesce when signaled by the everRun software.
- When you install the QEMU guest agent, you may need to restart your VMs. If your VMs are in use, schedule a maintenance period for the installation.

Related Topics

["Configuring Linux-based Virtual Machines" on page 221](#)

["Creating and Migrating Virtual Machines" on page 141](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

Managing the Operation of a Virtual Machine

Manage the operation of a virtual machine as described in:

- ["Starting a Virtual Machine" on page 225](#)
- ["Shutting Down a Virtual Machine" on page 225](#)
- ["Powering Off a Virtual Machine" on page 226](#)
- ["Opening a Virtual Machine Console Session" on page 227](#)
- ["Renaming a Virtual Machine" on page 228](#)
- ["Removing a Virtual Machine" on page 229](#)

For additional information configuration and troubleshooting information, see ["Advanced Topics \(Virtual Machines\)" on page 263](#).

Starting a Virtual Machine

Start a virtual machine (VM) to boot the VM's guest operating system.

To start a virtual machine

1. On the **Virtual Machines** page, select a VM.
2. Click **Start**.

Related Topics

["Shutting Down a Virtual Machine" on page 225](#)

["Powering Off a Virtual Machine" on page 226](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

Shutting Down a Virtual Machine

Shut down a virtual machine (VM) to begin an orderly shutdown of the guest operating system.



Note: You can shut down a VM with guest operating system commands. Some guests allow (or can be configured to allow) you to shut down a VM using the everRun Availability Console.

To shut down a VM in everRun Availability Console

1. On the **Virtual Machines** page, select a VM.
2. Click **Shutdown**.

A warning message appears, asking you to confirm the shutdown. Click **Yes** to shutdown or **No** to discontinue the shutdown.

If the VM is not responding, you can also **Power Off** the VM to stop it without properly shutting down the guest operating system.

Shutting down a VM 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:

- For any guest, verify that the power button action is set to shut down the guest operating system and not to hibernate it. If you click **Shutdown** in the everRun Availability Console for a guest that is set to hibernate, the VM remains in a **stopping** state and never properly shuts down.
- On some guests, the power button does not shut down the system unless a user is logged on to the operating system. You may be able to update security settings to enable the power button even in the absence of a login session.
- On some minimal server versions of Ubuntu 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 225](#)

["Powering Off a Virtual Machine" on page 226](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

Powering Off a Virtual Machine

Power off a virtual machine (VM) to stop it without properly shutting down guest operating system.



Caution: Use the **Power Off** command only if the **Shutdown** command or guest operating system commands fail. Powering off a VM is similar to pulling the power cord, which may result in data loss.

To power off a virtual machine

1. On the **Virtual Machines** page, select a VM.
2. Click **Power Off**.

Related Topics

["Starting a Virtual Machine" on page 225](#)

["Shutting Down a Virtual Machine" on page 225](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

["Advanced Topics \(Virtual Machines\)" on page 263](#)

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. On the **Virtual Machines** page, select a VM.
2. Ensure that the VM is in a running state.
3. Click **Console**.

Troubleshooting

To resolve an issue where the VM console window does not open

Ask your network administrator to open ports 6900-6999 (inclusive).

To resolve an issue where the VM console window is blank

Verify that the VM is powered on and not in the process of booting. Also, click in the console window and press any key to deactivate the screen saver.

To resolve an issue where more than one VM console window is displayed and they are behaving erratically

Close all console windows and open only one console window.

To resolve an issue where the VM console window hangs on the everRun system

For Ubuntu-based VMs, the VM console hangs in the everRun Availability Console if you do not properly set the `gfxmode` parameter. In the guest operating system, edit the `/boot/grub/grub.cfg` file and change the `gfxmode` parameter to `text` (for example, `set gfxmode=text`).

If the console hangs before you can set the parameter, do the following:

1. Restart the VM in the everRun Availability Console.
2. At the GRUB menu, press `e` to edit the grub command.
3. On the next screen, on the `gfxmode` line, change `$linux_gfx_mode` to `text` so the line reads:

```
gfxmode text
```

4. Press **Ctrl-x** or **F10** to boot the guest operating system.
5. To update the setting so it persists for each boot cycle, edit the `/boot/grub/grub.cfg` file and change the `gfxmode` parameter to `text` so the line reads:

```
set gfxmode=text
```

6. Save the `/boot/grub/grub.cfg` file.

To change the terminal type in a Linux-based VM if the console screen is unreadable

By default, the Linux operating system sets the `TERM` variable to `vt100-nav`, which is not properly supported by the `vncterm` program, the basis for the VM console in 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 225](#)

["Shutting Down a Virtual Machine" on page 225](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

Renaming a Virtual Machine

Rename a virtual machine (VM) to change its name as it appears on the **Virtual Machines** page.

If you need to change the host name of the guest operating system running in a VM, use guest operating system tools.



Prerequisite: To rename a VM, you must shut it down.

To rename a virtual machine

1. On the **Virtual Machines** page, select a VM.
2. Click **Shutdown** and wait for the VM to shut down.
3. Double-click the name of the VM.
4. Type the new name and press **Enter**.

Related Topics

["Removing a Virtual Machine" on page 229](#)

["Creating and Migrating Virtual Machines" on page 141](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

Removing a Virtual Machine

Remove a virtual machine (VM) to permanently delete it and optionally delete associated volumes from the everRun system.

Notes:



- When you remove a VM, any snapshots associated with the VM and the volume containers in which the snapshots are stored remain on the everRun system. To remove a VM snapshot and all of its associated volume snapshots, see ["Removing a Snapshot" on page 262](#).
- When all volume and volume snapshot contents have been removed from a volume container, the system automatically removes the container from the system, which frees up space in the storage group.



Prerequisite: Both PMs of the everRun system must be online to properly remove a VM. 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 remove a virtual machine

1. On the **Virtual Machines** page, select a VM.
2. Click **Shutdown**.
3. When the VM has stopped, click **Remove**.
4. In the **Remove Virtual Machine** dialog box, activate the check box next to volumes that you want to delete. Clear the check box for volumes to save as archives or save for attachment to another VM.



Caution: Make sure that you select the correct VM and volumes for removal. When you click **Delete VM**, these items are permanently removed.

5. Click **Delete VM** to permanently delete the VM and any selected volumes.

Related Topics

["Renaming a Virtual Machine" on page 228](#)

["Creating and Migrating Virtual Machines" on page 141](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

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 231](#)

To reconfigure virtual machine volumes, see the following task-specific topics:

- ["Creating a Volume in a Virtual Machine" on page 234](#)
- ["Attaching a Volume to a Virtual Machine" on page 235](#)

- ["Detaching a Volume from a Virtual Machine" on page 237](#)
- ["Removing a Volume from a Virtual Machine" on page 238](#)
- ["Expanding a Volume Container on the everRun System" on page 241](#)
- ["Expanding a Volume on the everRun System" on page 242](#)

To recover virtual machine resources, freeing space for new volumes or virtual CDs, see:

- ["Recovering Virtual Machine Resources" on page 243](#)

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 136](#). For more information about storage resources, see ["Planning Virtual Machine Storage" on page 139](#).
- 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 99](#)).
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, and Protection** 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:

- **Fault Tolerant (FT)**
- **High Availability (HA)**

For information about these levels of protection, see ["Creating a New Virtual Machine" on page 143](#) and ["Modes of Operation" on page 11](#).

c. Click **Next**.

5. On the **vCPUs and Memory** page:

a. Specify the number of **vCPUs** and the amount of **Memory** to assign to the VM. For more information, see ["Planning Virtual Machine vCPUs" on page 136](#) and ["Planning Virtual Machine Memory" on page 138](#).

b. If applicable, select the **Keyboard Layout in Console** to set the keyboard type for your country.

c. Click **Next**.

6. On the **Volumes** page, you can:

Notes:



- You cannot modify the VM boot volume, only data volumes
- To expand a volume container, see ["Expanding a Volume Container on the ever-Run System" on page 241](#).

- Click **Detach** to disconnect a volume from a VM and keep it for future use.
- Click **Delete** to permanently remove a volume from the everRunsystem.
- Select an unattached volume from a pulldown menu (if displayed) and click **Attach**.
- Add a new volume:
 - a. Type the **Name** of the boot volume.
 - b. Type the **Container Size** and **Volume Size** of the volume in gigabytes (GB). The container size is the total size for the volume including extra space to store snapshots. The volume size is the portion of the container that is available to the guest operating system. For more information about allocating storage, see ["Sizing Volume](#)

[Containers](#) on page 16 and [Planning Virtual Machine Storage](#) on page 139.

- c. Select the **Storage Group** for the volume, and, if applicable, select the **Sector Size**.

Ensure that you select a storage group that supports the sector size of the volume (see [Planning Virtual Machine Storage](#) on page 139). Note that the boot volume must have a sector size of 512 B. You can select the sector size, either 4K or 512B, only for data disks.

- d. If applicable, click **Add New Volume** to create a new data volume. (If the button is not visible, scroll down to the bottom of the wizard page.) Specify the parameters for the new volume.
- e. If applicable, click **Attach** to connect a volume to a VM.

To continue, click **Next**.

7. On the **Networks** page, activate the check box for each shared network that you want to attach to the VM.

For each shared network that you attach, you can also optionally:

- Set a custom MAC address (for details, see [Assigning a Specific MAC Address to a Virtual Machine](#) on page 263).
- Set the **State** to **Enabled** or **Disabled**, which allows you to allow or block network traffic to the selected network.

For more information, see [Planning Virtual Machine Networks](#) on page 141. 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 from 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 230](#)

["Planning Virtual Machine Resources" on page 136](#)

["Managing Virtual Machines" on page 135](#)

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 235](#).)



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 99](#)).
2. Select a VM and click **Shutdown**.
3. When the VM has stopped, click **Config** to display the **Reprovision Virtual Machine** wizard.
4. Click **Next** on each wizard page until the **Volumes** page is displayed. (If applicable, see ["Reprovisioning Virtual Machine Resources" on page 231](#) to configure additional VM resources.)
5. On the **Volumes** page, click **Add a new volume**. (If the button is not visible, scroll down to the bottom of the wizard page.)
6. Under **To Be Created**, do the following:
 - a. Type the **Name** of the volume as it will appear in the everRun Availability Console.
 - b. Type the **Container Size** and **Volume Size** of the volume to create in gigabytes (GB). The container size is the total size for the volume including extra space to store snapshots. The volume size is the portion of the container that is available to the guest operating system. For more information about allocating storage, see ["Sizing Volume Containers" on page 16](#) and ["Planning Virtual Machine Storage" on page 139](#).

- c. Select the **Disk Image** format:
 - **RAW** – raw disk format
 - **QCOW2** – QEMU Copy On Write (QCOW2) format, which supports snapshots
- d. Select the **Storage Group** in which to create the volume, and, if applicable, select the **Sector Size**.

Ensure that you select a storage group that supports the sector size of the volume that you want to create (see ["Planning Virtual Machine Storage" on page 139](#)). Note that the boot volume must have a sector size of 512 B. You can select the sector size, either 4K or 512B, only for data disks.

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 218](#)
 - ["Creating and Initializing a Disk \(Linux-based VMs\)" on page 222](#)

Related Topics

["Detaching a Volume from a Virtual Machine" on page 237](#)

["Removing a Volume from a Virtual Machine" on page 238](#)

["Managing Virtual Machine Resources" on page 230](#)

["Planning Virtual Machine Resources" on page 136](#)

["Managing Virtual Machines" on page 135](#)

Attaching a Volume to a Virtual Machine

Attach a volume to connect a currently unused volume to a virtual machine.



Note: If you attach a boot volume to a VM that already has a boot volume, the newly added volume is attached as a data volume. You might want to attach a volume in this manner to diagnose a boot problem or data corruption in another VM's boot volume. After using guest operating system tools to resolve the issue, detach the volume and reattach it to its original VM.



Prerequisite: Before attaching a volume to a virtual machine, you must shut down the virtual machine.

To attach a volume to a virtual machine

1. Ensure that the volume you want to attach is not in use by another virtual machine; otherwise, you cannot attach it. Open the **Volumes** page, locate the volume, and ensure that the value in the **Used By** column is **None**.
2. Open the **Virtual Machines** page (see ["The Virtual Machines Page" on page 99](#)).
3. Select a VM and click **Shutdown**.
4. When the VM has stopped, click **Config** to display the **Reprovision Virtual Machine** wizard.
5. Click **Next** on each wizard page until the **Volumes** page is displayed. (If applicable, see ["Reprovisioning Virtual Machine Resources" on page 231](#) 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 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 234](#)

["Detaching a Volume from a Virtual Machine" on page 237](#)

["Removing a Volume from a Virtual Machine" on page 238](#)

["Managing Virtual Machine Resources" on page 230](#)

["Planning Virtual Machine Resources" on page 136](#)

["Managing Virtual Machines" on page 135](#)

Detaching a Volume from a Virtual Machine

Detach a volume to disconnect it from a virtual machine and keep it for future use, or attach it to another virtual machine as described in ["Attaching a Volume to a Virtual Machine" on page 235](#). (You can also permanently delete the volume from the everRun system, as described in ["Removing a Volume from a Virtual Machine" on page 238](#).)

Notes:

- When you detach a volume from a VM, both the volume and its volume container exist separately from the VM. They remain on the system even if you remove the VM.
- If you decide to remove the volume, and you also want to remove its volume container to reclaim space in the storage group, you must remove any snapshots stored in the volume container; otherwise, the volume container remains on the system. For more information, see ["Removing a Volume from a Virtual Machine" on page 238](#).
- If you detach a boot volume from a VM, you cannot boot the VM; however, you might want to detach the boot volume to diagnose a boot problem or data corruption in the volume. You can temporarily attach the boot volume to another VM as a data volume, as described in ["Attaching a Volume to a Virtual Machine" on page 235](#). After using guest operating system tools to resolve the issue, detach the volume and reattach it to its original VM.



Prerequisite: Before detaching a volume from a virtual machine, you must shut down the virtual machine.

To detach a volume from a virtual machine

1. Open the **Virtual Machines** page (see ["The Virtual Machines Page" on page 99](#)).
2. Select a VM and click **Shutdown**.
3. When the VM has stopped, click **Config** to display the **Reprovision Virtual Machine** wizard.

4. Click **Next** on each wizard page until the **Volumes** page is displayed. (If applicable, see ["Reprovisioning Virtual Machine Resources" on page 231](#) to configure additional VM resources.)
5. On the **Volumes** page, locate the volume to detach. (If the volume is not visible, scroll down on the wizard page.)
6. Click **Detach** beside the volume name to mark the volume for detachment.



Caution: Be careful to mark the correct volume to detach, avoiding any volumes that are currently in use.

7. Click **Next** on each wizard page until the **Configuration Summary** page is displayed. Verify the configuration changes.
8. Click **Finish** to detach the selected volume.

Related Topics

["Attaching a Volume to a Virtual Machine" on page 235](#)

["Removing a Volume from a Virtual Machine" on page 238](#)

["Managing Virtual Machine Resources" on page 230](#)

["Planning Virtual Machine Resources" on page 136](#)

["Managing Virtual Machines" on page 135](#)

Removing a Volume from a Virtual Machine

Remove a virtual machine (VM) volume to permanently delete it from the everRun system. (You can also detach a volume from the VM but keep it for future use, as described in ["Detaching a Volume from a Virtual Machine" on page 237.](#))

Notes:

- If you remove a volume, and you also want to remove its volume container to reclaim space in the storage group, you must remove any volume snapshots stored in the volume container; otherwise, the container remains on the system. To remove a VM snapshot and all of its associated volume snapshots, see ["Removing a Snapshot" on page 262](#).
- When all volume and volume snapshot contents have been removed from a volume container, the system automatically removes the container from the system, which frees up space in the storage group.



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 99](#)).
2. Select a VM and click **Shutdown**.
3. When the VM has stopped, click **Config** to display the **Reprovision Virtual Machine** wizard.
4. Click **Next** on each wizard page until the **Volumes** page is displayed. (If applicable, see ["Reprovisioning Virtual Machine Resources" on page 231](#) to configure additional VM resources.)
5. On the **Volumes** page, locate the volume to delete. (If the volume is not visible, scroll down on the wizard page.)
6. Click **Delete** beside the volume name to mark the volume for deletion.



Caution: Be careful to mark the correct volume to remove, avoiding any volumes that are currently in use.

7. Click **Next** on each wizard page until the **Configuration Summary** page is displayed. Verify the configuration changes.
8. Click **Finish** to permanently delete the selected volume.

To remove an unattached volume



Caution: Before removing a volume, ensure that it is no longer needed by other administrators.

1. Open the **Volumes** page.
2. Select an unattached volume. (The **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 237](#)

["Attaching a Volume to a Virtual Machine" on page 235](#)

["Managing Virtual Machine Resources" on page 230](#)

["Planning Virtual Machine Resources" on page 136](#)

["Managing Virtual Machines" on page 135](#)

Renaming a Volume on the everRun System

Rename a volume on the everRun 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 234](#)

["Detaching a Volume from a Virtual Machine" on page 237](#)

["Removing a Volume from a Virtual Machine" on page 238](#)

["Managing Virtual Machine Resources" on page 230](#)

["Planning Virtual Machine Resources" on page 136](#)

["Managing Virtual Machines" on page 135](#)

Expanding a Volume Container on the everRun System

Expand a virtual machine (VM) volume container to allocate more space in the container for snapshots or for the guest operating system volume. (To expand the portion of a volume container that is available to the guest operating system, see ["Expanding a Volume on the everRun System" on page 242.](#))

You can expand the volume container, but you cannot reduce the size of a container. Use the following procedure to expand a volume container whether the VM is running or stopped. To estimate the amount of storage to allocate to a volume container, see ["Sizing Volume Containers" on page 16.](#)



Prerequisite: Ensure that both PMs of the everRun system are online; otherwise, the system cannot properly expand a volume container.

To expand a volume container

1. On the **Physical Machines** page (see ["The Physical Machines Page" on page 95](#)), 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 **Volumes** page (see ["The Volumes Page" on page 104](#)), click **Expand Container**.
3. Next to **Expand By**, type the amount of storage space to add to the volume container (in gigabytes (GB)). When you type the number, the dialog box displays the **Expanded Container Size** that will result if you complete the operation.



Note: Consider the **Expand By** entry carefully, because after expanding a container, you cannot undo the change or reduce the size of the volume container; you can only expand the volume further.

4. Click **Expand Container** to commit the change and expand the container. The dialog box displays the expansion progress and automatically closes when the operation is complete.

Related Topics

["Expanding a Volume on the everRun System" on page 242](#)

["Creating a Volume in a Virtual Machine" on page 234](#)

["Detaching a Volume from a Virtual Machine" on page 237](#)

["Removing a Volume from a Virtual Machine" on page 238](#)

["Managing Virtual Machine Resources" on page 230](#)

["Planning Virtual Machine Resources" on page 136](#)

["Managing Virtual Machines" on page 135](#)

Expanding a Volume on the everRun System

Expand a virtual machine (VM) volume to allocate more space for programs and data in the guest operating system. Before expanding a VM volume, you may also need to expand its volume container, as described in ["Expanding a Volume Container on the everRun System" on page 241](#), to ensure that the volume container has enough space to expand the volume and store snapshots.

You can expand a volume, but you cannot reduce the size of a volume. Use the following procedure to expand a volume only when the VM is stopped.

Prerequisites:



- You must shut down the VM before expanding a volume that it contains.
- Ensure that both PMs of the everRun system are online; otherwise, the system cannot properly expand a volume.

To expand a volume container

1. On the **Physical Machines** page (see ["The Physical Machines Page" on page 95](#)), verify that both PMs are in the **running** state and that neither PM is in maintenance mode or in the process of synchronizing.
2. If applicable, expand the volume container for the volume as described ["Expanding a Volume Container on the everRun System" on page 241](#). The volume container must have at least as much available space as the amount of space you want to add to the volume. If you also take snapshots, additional space is needed.
3. On the **Virtual Machines** page (see ["The Virtual Machines Page" on page 99](#)), select the VM that contains the volume that you want to expand. Ensure that the VM is **stopped**.
4. In the bottom pane, click the **Volumes** tab and select the volume that you want to expand. In the **Action** column, click **Expand Volume**.

5. Next to **Expand By**, type the amount of storage space to add to the volume (in gigabytes (GB)). When you type the number, the dialog box displays the **Expanded Volume Size** that will result if you complete the operation.



Note: Consider the **Expand By** entry carefully, because after expanding a volume, you cannot undo the change or reduce the size of the volume; you can only expand the volume further.

6. Click **Expand Volume** to commit the change and expand the volume. The dialog box displays the expansion progress and automatically closes when the operation is complete.

Related Topics

["Expanding a Volume Container on the everRun System" on page 241](#)

["Creating a Volume in a Virtual Machine" on page 234](#)

["Detaching a Volume from a Virtual Machine" on page 237](#)

["Removing a Volume from a Virtual Machine" on page 238](#)

["Managing Virtual Machine Resources" on page 230](#)

["Planning Virtual Machine Resources" on page 136](#)

["Managing Virtual Machines" on page 135](#)

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 229](#)
- ["Removing a Volume from a Virtual Machine" on page 238](#)
- ["Removing a Virtual CD" on page 248](#)

You can also remove unused snapshots from a VM to free space for new snapshots on an existing volume, but doing so does not recover storage space for new volumes or VCDs:

- ["Removing a Snapshot" on page 262](#)

Related Topics

["Managing Virtual Machine Resources" on page 230](#)

["Planning Virtual Machine Resources" on page 136](#)

["Managing Virtual Machines" on page 135](#)

Managing Virtual CDs

Create and manage virtual CDs (VCDs) to make software installation media available to the virtual machines on your everRun 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, as described in ["Creating a Virtual CD" on page 244](#).

After you create a VCD, you can boot from it to install a Windows or Linux guest operating system, or start a VM from a bootable recovery VCD. You can also insert a VCD into a running VM to install software applications.



Caution: When you insert a VCD into a running, fault-tolerant (FT) VM, it prevents the everRun software from migrating the VM to a different physical machine in the event of a failure. To restore fault-tolerant operation, unmount and eject the VCD as soon as you finish using it.

You manage VCDs as described in:

- ["Creating a Virtual CD" on page 244](#)
- ["Inserting a Virtual CD" on page 246](#)
- ["Ejecting a Virtual CD" on page 246](#)
- ["Booting from a Virtual CD" on page 247](#)
- ["Renaming a Virtual CD" on page 248](#)
- ["Removing a Virtual CD" on page 248](#)

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 upload or 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 247](#)) to install a guest operating system or start a VM from a bootable recovery VCD. You can also insert a VCD into a running VM (see ["Inserting a Virtual CD" on page 246](#)) to install software applications.

Notes:



1. 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.
2. If you create a bootable VCD for installation, it must be a single CD or DVD. Multiple CDs or DVDs are not supported.

To create a VCD

1. If necessary, create ISO files of any physical media for which you will create VCDs.
2. Open the **Virtual CDs** page in the 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 system running the everRun Availability Console. Click **Browse**, select the ISO file on your system, and click **Open**.
 - **Copy CD ISO from network source** copies the file from a Web URL. Specify the URL of the ISO file.
7. 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 source and size of the image.

You can determine the status of a VCD by checking the **State** column on the **Virtual CDs** page:

- A syncing icon () indicates that the VCD is still being created.
- A broken icon () indicates that the VCD creation failed. Remove the VCD and try creating it

again.

- A normal icon (✔) indicates that the transfer is complete and that the VCD is ready to use.

Related Topics

["Inserting a Virtual CD" on page 246](#)

["Ejecting a Virtual CD" on page 246](#)

["Managing Virtual CDs" on page 244](#)

["Creating and Migrating Virtual Machines" on page 141](#)

Inserting a Virtual CD

Insert a virtual CD (VCD) in a virtual machine (VM) to access installation media when installing applications in a guest operating system. (To boot a virtual machine from a VCD, see ["Booting from a Virtual CD" on page 247.](#))



Caution: When you insert a VCD into a running, fault-tolerant (FT) VM, it prevents the everRun software from migrating the VM to a different physical machine in the event of a failure. To restore fault-tolerant operation, unmount and eject the VCD as soon as you finish using it.

To connect a VCD to a VM

1. If necessary, create a VCD (see ["Creating a Virtual CD" on page 244](#)) for the software installation media you need to access.
2. On the **Virtual Machines** page, select a VM.
3. In the bottom pane, click the **CD Drives** tab.
4. On the **CD Drives** tab, select a VCD from the pulldown menu.
5. Click **Insert a CD** to connect the VCD to the VM.

Related Topics

["Creating a Virtual CD" on page 244](#)

["Ejecting a Virtual CD" on page 246](#)

["Booting from a Virtual CD" on page 247](#)

["Managing Virtual CDs" on page 244](#)

Ejecting a Virtual CD

Eject a virtual CD (VCD) to disconnect it from a virtual machine (VM). Ejecting a VCD allows you to insert another VCD into the VM. It also makes the VCD available for inserting into another VM.

To eject a VCD from a VM

1. Unmount the VCD from the guest operating system to ensure that it is not in use.
2. On the **Virtual Machines** page, select a VM.
3. Click the **CD Drives** tab in the lower frame.
4. On the **CD Drives** tab, click **Eject CD**.

Related Topics

["Creating a Virtual CD" on page 244](#)

["Inserting a Virtual CD" on page 246](#)

["Booting from a Virtual CD" on page 247](#)

["Managing Virtual CDs" on page 244](#)

Booting from a Virtual CD

Boot a virtual machine from a virtual CD (VCD) to install a guest operating system or to perform maintenance.

Before booting from a VCD, you must shut down the virtual machine.

To boot a virtual machine from a VCD

1. If necessary, create a VCD from a bootable CD/DVD (see ["Creating a Virtual CD" on page 244](#)).
2. On the **Virtual Machines** page, select a virtual machine.
3. If the virtual machine is running, click **Shutdown**.
4. When the virtual machine status shows **stopped**, click **Boot from CD** in the lower pane.
5. Select the bootable VCD, then click **Boot**.



Note: A Windows-based virtual machine booted from a VCD boots as a hardware virtual machine (HVM), and it can access only the first three disk volumes.

Related Topics

["Creating a Virtual CD" on page 244](#)

["Inserting a Virtual CD" on page 246](#)

["Ejecting a Virtual CD" on page 246](#)

["Managing Virtual CDs" on page 244](#)

["Creating and Migrating Virtual Machines" on page 141](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

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 248](#)

["Inserting a Virtual CD" on page 246](#)

["Ejecting a Virtual CD" on page 246](#)

["Creating a Virtual CD" on page 244](#)

["Managing Virtual CDs" on page 244](#)

Removing a Virtual CD

Remove a virtual CD (VCD) to permanently delete it from the everRun 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** column displays **Yes** for the VCD. If the value is **No**, the VCD is currently in use.
4. Select the VCD and click **Remove** in the lower panel.

Related Topics

["Renaming a Virtual CD" on page 248](#)

["Inserting a Virtual CD" on page 246](#)

["Ejecting a Virtual CD" on page 246](#)

["Creating a Virtual CD" on page 244](#)

["Managing Virtual CDs" on page 244](#)

Managing Snapshots

Snapshots allow you to save an image of a virtual machine (VM) or of selected volumes on a VM at a particular point in time. You can use a snapshot to create a new VM on the same everRun system or you can export the snapshot to files on a network share for use on another everRun system.



Caution: Creating a snapshot results in the conversion of any RAW format volume to QCOW2 format, which may have performance implications for your system. You cannot convert the volumes back to RAW format; therefore, if you have a particular requirement for RAW format volumes, avoid using the snapshot feature.



Notes:

- You cannot revert the state of a VM to a snapshot; however, you can create a new VM from a snapshot or export files that you use to restore or duplicate the original VM.
- When you create a snapshot, all volumes are selected, by default. You can, though, change the selection of individual volumes.
- The boot volume is required for all snapshots.

You manage snapshots as described in:

- ["Creating a Snapshot" on page 250](#)
- ["Creating a Virtual Machine from a Snapshot" on page 254](#)
- ["Exporting a Snapshot" on page 256](#)
- ["Removing a Snapshot" on page 262](#)

The everRun system's ability to take snapshots is enabled, by default. To disable or to re-enable the system's ability to take snapshots, see ["Disabling and Enabling Snapshots" on page 81](#).

To view the snapshots that you have created in the everRun Availability Console:

- Open the **Snapshots** page (see ["The Snapshots Page" on page 104](#))
- On the **Virtual Machines** page (see ["The Virtual Machines Page" on page 99](#)), click a VM and click the **Snapshots** tab.

When you create a VM snapshot, the everRun system saves a snapshot image that includes any data that has changed in the VM since the previous snapshot, or, if no snapshots exist, since you originally created the VM. Because each snapshot contains only the changed data, some snapshots may take a small amount of storage space, and other snapshots may take more space depending on the level of VM activity and the amount of time that has passed since the previous snapshot.

Because snapshots are stored in the volume containers for each volume, ensure that you reserve enough storage space in the volume container for each volume you want to include in your VM snapshots, as described in ["Sizing Volume Containers" on page 16](#). You can also remove older or obsolete snapshots to recover storage space.

You can create a snapshot of a VM whether the VM is running or shut down; however, if you want to create an *application-consistent* snapshot, where supported applications *quiesce* or freeze their operations to ensure data consistency, you must prepare your guest operating system as described in one of the following topics:

- ["Installing the QEMU Guest Agent for Application-Consistent Snapshots \(Windows-based VMs\)" on page 219](#)
- ["Installing the QEMU Guest Agent for Application-Consistent Snapshots \(Linux-based VMs\)" on page 223](#)

Related Topics

["Creating and Migrating Virtual Machines" on page 141](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

Creating a Snapshot

Create a snapshot to save an image of a virtual machine (VM) or of selected volumes on a VM at a particular point in time. You can use a snapshot to create a new VM on the same everRun system or you can export the snapshot to files on a network share for use on another everRun system. By default, the everRun system's ability to take snapshots is enabled. To disable or to re-enable the system's ability to take snapshots, see ["Disabling and Enabling Snapshots" on page 81](#). For an overview of snapshots, see ["Managing Snapshots" on page 249](#).

You can create a snapshot of a VM whether the VM is running or shut down; however, if you want to create an *application-consistent* snapshot, where supported applications *quiesce* or freeze their operations to ensure data consistency, you must prepare your guest operating system as described in one of the following topics:

- ["Installing the QEMU Guest Agent for Application-Consistent Snapshots \(Windows-based VMs\)" on page 219](#)
- ["Installing the QEMU Guest Agent for Application-Consistent Snapshots \(Linux-based VMs\)" on page 223](#)

The number of snapshots you can create depends on the amount of storage space you have allocated in the volume container for each VM volume, as described in ["Sizing Volume Containers" on page 16](#). If necessary, you can expand a volume container as described in ["Expanding a Volume Container on the everRun System" on page 241](#).



Caution: Creating a snapshot results in the conversion of any RAW format volume to QCOW2 format, which may have performance implications for your system. You cannot convert the volumes back to RAW format; therefore, if you have a particular requirement for RAW format volumes, avoid using the snapshot feature.

Notes:



- For Linux-based VMs, if you want to create a snapshot of the VM to export to another system, 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 may use different device names on another 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.
- If you want to shut down the source VM while creating a snapshot, consider scheduling a planned maintenance period for this process.
- When you create a snapshot, all volumes are selected, by default. You can, though, change the selection of individual volumes.
- The boot volume is required for all snapshots.
- If you want to use a snapshot to duplicate a VM, and you will continue to use the source VM after the export, remember to set a different MAC address and IP address for the VM when you import it on the target system.
- If the everRun system switches from the primary PM to the secondary PM during the snapshot, the snapshot fails. This does not affect the continuous uptime of your system, but the snapshot is automatically deleted, and you need to start a new snapshot.



Prerequisite: Both PMs of the everRun system must be online for the snapshot process to function properly. If only one PM is online, the snapshot is written only to the online PM, and that same PM must be primary if you export the snapshot later.

To prepare for creating a snapshot (Windows-based VMs only)

1. If you want to create an application-consistent snapshot, ensure that the QEMU guest agent is installed and running, as described in ["Installing the QEMU Guest Agent for Application-Consistent Snapshots \(Windows-based VMs\)"](#) on page 219.
2. Ensure that all volumes are labeled accurately, as summarized in ["Managing Windows Drive Labels"](#) on page 214.

3. Run the Windows System Preparation Tool (`Sysprep`) if you need to prepare the guest operating system for redeployment.

To prepare for creating a snapshot (Linux-based VMs only)

If you want to create an application-consistent snapshot, ensure that the QEMU guest agent is installed and running, as described in ["Installing the QEMU Guest Agent for Application-Consistent Snapshots \(Linux-based VMs\)" on page 223](#).

To create a snapshot

1. Log on to the everRun system with the everRun Availability Console.
2. On the **Physical Machines** page (see ["The Physical Machines Page" on page 95](#)), 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, select a VM.
4. With the VM selected, click the **Snapshot** button in the bottom pane.
5. In the **Snapshot Virtual Machine** dialog box, all volumes (under **Volumes to Capture**) are selected, by default. Deactivate the check box beside volumes that you do not want to capture in the snapshot. The boot volume is required for all snapshots.

Optionally, type a **Snapshot Name** and **Description** for the snapshot. The default **Snapshot Name** for each new snapshot is the name of the VM, but you can type a more descriptive name. (The snapshot name does not need to be unique.)

6. Click **Create Snapshot**. The snapshot begins, and the dialog box closes automatically.

Creating a snapshot typically takes a few seconds, but it may take longer depending on the level of VM activity and the amount of time that has passed since the previous snapshot. You can determine the status of a snapshot by checking the **State** column on the **Snapshots** page:

- A broken icon (✘) indicates that a snapshot is still in progress, or that is written to only one node in the everRun system.
- A normal icon (✔) indicates that a snapshot is complete.

If you want to use a completed snapshot to create a new VM, see ["Creating a Virtual Machine from a Snapshot" on page 254](#). If you want to export a completed snapshot, see ["Exporting a Snapshot" on page 256](#).

Related Topics

["Managing Snapshots" on page 249](#)

["Creating and Migrating Virtual Machines" on page 141](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

Creating a Virtual Machine from a Snapshot

Create a virtual machine (VM) from a snapshot if you want to use a snapshot on your everRun system as the source for a new VM. (For additional methods of creating or migrating VMs, see ["Creating and Migrating Virtual Machines" on page 141](#). For an overview of snapshots, see ["Managing Snapshots" on page 249](#).)

To create a VM from a snapshot, open the **Snapshots** page of everRun Availability Console, select a snapshot, and click **Create VM**. A wizard steps you through the process of creating the VM, as described in the following procedure.

Notes:

- When you create a snapshot that you intend to use as a source for a new VM, you must follow steps to prepare the guest operating system; otherwise, the VM image that you create may not function as expected. For details, see ["Creating a Snapshot" on page 250](#).
- When you create a VM from a snapshot, the original container size for each volume you include is not preserved. For example, if your source VM has a 20 GB boot volume in a 40 GB volume container, the new VM will have a 20 GB boot volume in a 20 GB volume container. If necessary, you can expand the volume containers for the new VM as described in ["Expanding a Volume Container on the everRun System" on page 241](#).
- To prevent conflicts with the source VM, the create VM wizard automatically assigns a new MAC address to each network interface in the new VM; however, you may need to manually update any IP addresses and host names.



Prerequisite: Ensure that both PMs of the everRun system are online; otherwise, the system cannot properly create the VM.

To create a new VM from a snapshot

1. Log on to the everRun system with the everRun Availability Console.
2. On the **Physical Machines** page (see ["The Physical Machines Page" on page 95](#)), verify that both PMs are in the **running** state and that neither PM is in maintenance mode or in the process of synchronizing.
3. If you have not already done so, create a snapshot, as described in ["Creating a Snapshot" on page 250](#).
4. On the **Snapshots** page, select the snapshot to use as a source for the new VM.
Snapshots are usually in a normal state (✓) the **State** column. If a snapshot is broken (✗), it may indicate that one or more volumes in the snapshot are unavailable for the VM creation.
5. Click **Create VM**.
6. When prompted, type a new VM name and new volume names for the VM. The names must be unique on the everRun system.
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 system (the boot volume is required). Select the **Restore Data** box to import data for a volume from the snapshot.
 - **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.
8. Optionally, clear the check box for **Auto start Virtual Machine** if you need to reprovision the VM before starting it for the first time.
9. Click **Create VM**. When the process is complete, the wizard closes automatically.
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 231](#). Also, if you want to

allocate additional space in each volume container for snapshots, see ["Expanding a Volume Container on the everRun System" on page 241](#).

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

Related Topics

["Managing Snapshots" on page 249](#)

["Creating and Migrating Virtual Machines" on page 141](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

Exporting a Snapshot

Export a snapshot to transfer a virtual machine (VM) image from an everRun system to a network-mounted folder (that is, directory). Exporting a snapshot makes the VM image available for importing to another system or for importing back to the same everRun system to restore or duplicate the original VM. (For an overview of snapshots, see ["Managing Snapshots" on page 249](#). For additional VM migration/export methods, see ["Creating and Migrating Virtual Machines" on page 141](#).)

You prepare for exporting snapshots by creating a network-mounted folder to store an exported VM in your environment. Create a folder for either a Windows share (also known as a Common Internet File System (CIFS) share) or a Network File System (NFS) export. After creating the folder, you mount it in the host operating system of the everRun system, as described in this topic. When you initiate an export in the everRun Availability Console, the everRun system saves the VM image to the network-mounted folder as standard Open Virtualization Format (OVF) and Virtual Hard Disk (VHD) files.

Notes:

- When you create a snapshot that you intend to export, you must follow steps to prepare the guest operating system; otherwise, the VM image that you create may not function as expected. For details, see ["Creating a Snapshot" on page 250](#).
- When you export a snapshot, you export a fully coalesced snapshot of the VM from that point in time, not only the changed data. If you want to create differential backups of a VM, use a third-party backup solution.
- When you export a snapshot to import a VM to another everRun system, the original container size for each volume you include is not preserved. For example, if your source VM has a 20 GB boot volume in a 40 GB volume container, the target VM will have a 20 GB boot volume in a 20 GB volume container. If necessary, you can expand the volume containers on the target everRun system as described in ["Expanding a Volume Container on the everRun System" on page 241](#).
- The time required for the export depends on the size and number of volumes in the source VM as well as network bandwidth. For example, transferring a VM with one 20 GB boot disk over a 1 Gb/s network may take about 30 minutes.
- If you will continue to use the source VM after the export, remember to set a different MAC address and IP address for the VM when you import it on the target system.
- If the everRun system switches from the primary PM to the secondary PM during an export, the export process fails. This does not affect the continuous uptime of your system. You can delete the partially exported files from your system running the everRun Availability Console and export the files again.



Prerequisite: Both PMs of the everRun system must be online for the export process to function properly. You can export a snapshot from a single-node system only if all of the volume snapshots that you select to include in the export are present on the primary node, as displayed in the **Export Snapshot** dialog box. In most cases, snapshots are replicated on both nodes, but a snapshot may be unavailable if a node was offline when the snapshot was taken.

**To create and mount an NFS export or Windows CIFS share**

Before you can export a snapshot, you must create and mount the folder to which the export(s) will

be transferred, as follows:

1. Create a network-mounted folder for a Windows/CIFS share or an NFS export in your environment where you can store VM exports.

Ensure that you have enough storage for the VMs that you want to export. Also, set full read/write permissions on the network-mounted folder to permit file transfers, or, for a Windows/CIFS share only, assign read/write permissions to a specific user on the system/domain that hosts the share. Record the location and settings of the NFS export or CIFS share, which you will specify in a later step.

2. Log on to the everRun system with the everRun Availability Console.
3. On the **Physical Machines** page, verify that both PMs are in the running state and that neither PM is in maintenance mode or in the process of synchronizing. See ["The Physical Machines Page" on page 95](#). Also, make note of which PM is the **primary** node. In the list in the top pane, the primary node is marked as **noden (primary)**.
4. Obtain the IP address of the primary node, if you do not already know it. For example, on the **Preferences** page, click **IP Configuration**. Click the **Noden IP** tab for the primary node and record the **IP Address** value.
5. Use a secure shell (`ssh`) utility to log on to the host operating system (host OS) of the primary node of the everRun system, where you will mount the folder. Log on as the `root` user.

The next step shows how to use the `ftxmnt` script to automatically mount the folder. The script should work in most cases, but you can also manually mount the folder by executing standard mount commands if needed.

Notes:

- If you manually mount the NFS export or Windows/CIFS share in the host OS of the everRun system, you must create the mount point at `/mnt/ft-export`, where the export process expects to find it. (If you use the `ftxmnt` script, it automatically creates this mount point.)
- If you want the mount to persist when you reboot the everRun system, manually add an entry to the `/etc/ftstab` file in the host OS of the everRun system. (The `ftxmnt` script does not modify the `/etc/ftstab` file.)
- You need to mount the NFS export or Windows/CIFS share only on the primary node of the everRun system, but you can optionally add the mount to the `/etc/ftstab` file on both nodes so it is always available even if the primary node changes.
- If you mount a Windows/CIFS share that requires a password, and the password changes after the mount, you need to unmount the folder and mount it again with the new password; otherwise, the export may fail unexpectedly.

6. To automatically mount the NFS export or Windows/CIFS share, execute the `ftxmnt` script and respond to the interactive prompts. For example, the following output shows how to mount a Windows/CIFS share (`\\192.168.0.111\ExportVMs`) that is accessible by a specific user account:

```
[root@node0 /]# ftxmnt
```

```
This script is meant to mount a Network Attached Storage
location to use for exporting everRun virtual machines.
Enter Ctrl-C to exit
```

```
Enter n if you are mounting an nfs share, enter w if you
are entering a windows share: w
```

```
What is the IP address or the computer name of the file
```

server?

192.168.0.111

What is the name of the share you wish to mount?

ExportVMs

Does this share require authentication? (y/n):

Respond **y** to the prompt above if the share requires authentication. When you respond **y**, the following prompts appear:

- What is your username?

Enter your *domain\username*.

If the share does not require authentication, type **n** and then respond to the following prompt:

If there is a guest account name, enter it here:

Type the guest account name, if it exists.

The following message appears:

```
Successfully mounted folder \\192.168.0.111\ExportVMs
at path /mnt/ft-export/
```

7. Switch to the `/mnt/ft-export` directory in the everRun host OS and create a file to verify that the export or share is present and that you have read/write permissions. For example:

```
# touch test
# ls
test
```

Confirm that the file also appears in the export or share on the remote system. If the file is not present or the everRun host OS displays an error, verify your mount settings and permissions.

8. Remove the test file.

```
# rm test
```

Later, you want to unmount the export or share after exporting your virtual machines, switch out of the `/mnt/ft-export` directory and execute the `umount` command, as follows:

```
# cd /
# umount /mnt/ft-export
```

To export a snapshot

1. Log on to the everRun system with the everRun Availability Console.
2. On the **Physical Machines** page (see ["The Physical Machines Page" on page 95](#)), verify that both PMs are in the **running** state and that neither PM is in maintenance mode or in the process of synchronizing.
3. If you have not already done so, create a snapshot, as described in ["Creating a Snapshot" on page 250](#).

4. On the **Snapshots** page, select the snapshot to export.

Snapshots are usually in a normal state (✓) the **State** column. If a snapshot is broken (✗), it may indicate that one or more volumes in the snapshot are unavailable for export. You can check volume availability in step 7.

5. Click **Export**.
6. In the **Export Snapshot** dialog box, type the path in **/mnt/ft-export** where you want the snapshot to be exported.

For example, if you want the export process to create a new directory named `ocean1` to store the OVF and VHD files, type `ocean1`. If you want the export process to create the `ocean1` directory inside an existing directory named `TestVMs`, type `TestVMs/ocean1`.

7. Review the list of **Captured Volumes** and select the volumes to include in the export.

In most cases, the dialog box shows that **All Captured Volumes are available for Export on node n**, the primary node. You can select any snapshots for export.

If one or more snapshots are unavailable on the primary node (usually because the node was offline when the snapshot was taken), the dialog box allows you to select only the available snapshots. If necessary, you can cancel the export, ensure that both nodes are in the **running** state, and create a new snapshot for export.

8. Click **Export Snapshot**. The export begins, and the dialog box closes automatically.

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

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

After the export process, if you want to import or restore the OVF and VHD files on an everRun system, see ["Importing an OVF File from an everRun 7.x System " on page 198](#)

Troubleshooting

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

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

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

Related Topics

["Managing Snapshots" on page 249](#)

["Creating and Migrating Virtual Machines" on page 141](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

Removing a Snapshot

Remove a snapshot to permanently delete it from the everRun system.

Notes:



- When you remove a VM snapshot, you also remove all of its associated volume snapshots, which frees up storage space in the volume containers that contain those volume snapshots.
- If you remove all volume and volume snapshot contents from a volume container, the system automatically removes the container from the system, which frees up space in the storage group.

To remove a snapshot

1. On the **Snapshots** page, select a snapshot to remove.
2. Click **Delete**.

3. A confirmation window appears. Click **Yes** to continue deleting the snapshot or click **No** to stop the deletion.

Related Topics

["Managing Snapshots" on page 249](#)

["Creating and Migrating Virtual Machines" on page 141](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

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 263](#)
- ["Selecting a Preferred PM for a Virtual Machine" on page 265](#)
- ["Changing the Protection Level for a Virtual Machine \(HA or FT\)" on page 265](#)
- ["Configuring the Boot Sequence for Virtual Machines" on page 266](#)
- ["Resetting MTBF for a Failed Virtual Machine" on page 267](#)
- ["Locating a Dump File in a Virtual Machine" on page 268](#)

To manage the operation of a virtual machine, see ["Managing the Operation of a Virtual Machine" on page 224](#).

Assigning a Specific MAC Address to a Virtual Machine

Assign a specific Media Access Control (MAC) address to a virtual machine (VM) if you need to override its default MAC address.

Warnings:



1. By default, the everRun software automatically assigns MAC addresses to the VMs. Do not override the default settings unless you have specific requirements (for example, to support software applications that are licensed on a MAC-address basis).
2. If you change the **Static System IP** address, any MAC addresses automatically assigned to the VMs will change when the VMs reboot, because the everRun software generates MAC addresses for the VMs based on the system IP address. To prevent changes to the MAC address for a VM, set a persistent MAC address as described in the following procedure. Contact your network administrator to generate a valid MAC address for your environment, and remember to update any firewall rules based on the new MAC address.



Prerequisite: Before modifying the MAC address for a virtual machine, you must shut down the VM.

To assign a specific MAC address to a VM

1. Open the **Virtual Machines** page (see ["The Virtual Machines Page" on page 99](#)).
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 231](#) 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 263](#)

["Managing Virtual Machine Resources" on page 230](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

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 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 263](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

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 225](#) for information about stopping a VM.)
2. In the bottom pane, click **Config** to open the **Reprovision Virtual Machine** wizard
3. On the **Name, Description and Protection** page, select the **HA** or **FT** button.
4. Continue clicking through the wizard pages to the last page. 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 263](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

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 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 263](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

Resetting MTBF for a Failed Virtual Machine

Reset the mean time between failure (MTBF) counter for a virtual machine to attempt to restart a failed virtual machine.

If a virtual machine's guest OS crashes, everRun 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.

Notes:



1. The **Reset Device** button is displayed only if the virtual machine falls below its MBTF threshold.
2. The **Clear MTBF** button is displayed only if the system software supporting a VM on one physical machine falls below its MBTF threshold.

To reset the MTBF counter for a virtual machine

1. On the **Virtual Machines** page, select a virtual machine.
2. Click **Reset Device**.

If the system software supporting a VM on one physical machine fails too often, perform the steps below to reset its MTBF counter.

To reset the MTBF counter for a VM on one physical machine

1. On the **Virtual Machines** page, select a virtual machine.
2. Click **Clear MTBF**.

Related Topics

["Advanced Topics \(Virtual Machines\)" on page 263](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

["Creating a Diagnostic File" on page 85](#)

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 263](#)

["Managing the Operation of a Virtual Machine" on page 224](#)

["Creating a Diagnostic File" on page 85](#)

8

Chapter 8: Maintaining Physical Machines

You can maintain PMs in an everRun 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 270.

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 270 for hot-swappable components such as network cables, fans, and power supplies
 - ["Adding or Replacing Components That Are Not Hot-Swappable"](#) on page 271 for components such as CPUs and memory or any other component that is not hot-swappable.
 - ["Adding a New NIC"](#) on page 272 for adding new network interface cards (NICs).
- To replace a PM or a failed motherboard, NIC, or RAID controller, see ["Replacing Physical Machines, Motherboards, NICs, or RAID Controllers"](#) on page 274.
- To replace a PM that is connected to external storage, see ["Replacing a Failed Physical Machine Connected to External Storage"](#) on page 277.
- To upgrade both PMs in a running system, see ["Upgrading Both Physical Machines In a Running System"](#) on page 278.

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 24](#) 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 125](#)

["Maintaining Physical Machines" on page 269](#)

["The everRun Availability Console" on page 58](#)

["Physical Machines and Virtual Machines" on page 9](#)

["The Physical Machines Page" on page 95](#)

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 270.

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 125](#)

["Maintaining Physical Machines" on page 269](#)

["The everRun Availability Console" on page 58](#)

["Physical Machines and Virtual Machines" on page 9](#)

["The Physical Machines Page" on page 95](#)

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 270.

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 125](#)

["Maintaining Physical Machines" on page 269](#)

["The everRun Availability Console" on page 58](#)

["Physical Machines and Virtual Machines" on page 9](#)

["The Physical Machines Page" on page 95](#)

Adding a New NIC

When adding new NICs, you must add NICs to both physical machines (PMs) and then cable the NICs to the appropriate switch on both sides in order to establish connectivity and to form one or more shared networks that you can then assign to VMs or use as A-links.



Prerequisite: Before you add a NIC, see ["Physical Machine Hardware Maintenance Restrictions" on page 270](#).

To add new NICs



Note: You can begin this procedure with **node0** or **node1** and then continue with the other node. The procedure below begins with **node0** for simplicity.

1. In the everRun Availability Console, click **Physical Machines** in the left-hand navigation panel.
2. Perform the following for **node0**:

- a. Select the appropriate node and then click **Work On**.
- b. After the node displays **running (in Maintenance)**, click **Shutdown** and then **OK**.
- c. Insert the new NIC in the desired slot.
- d. Press the power button to power on the node.

Wait for the PM to boot and for the everRun Availability Console to display **running** as the **Activity** state for the appropriate node under **Physical Machines**.

- e. Click **Finalize** for and then click **OK**, which exits the node from maintenance mode.
Wait while storage synchronization completes (🔄 disappears).

3. Perform Step 2 for **node1**.

In node1 insert the new NIC in the slot that corresponds to the slot where you inserted the new NIC in the PM that is node0 (Step c, above).

4. Connect network cables to the new NICs, as needed, and configure the new network as an A-Link or a business network. See ["Connecting Additional Networks" on page 55](#).
5. Reconfigure and start any VMs that need to use the new networks. See ["Managing Virtual Machines" on page 135](#).

Related Topics

["Maintenance Mode" on page 125](#)

["Maintaining Physical Machines" on page 269](#)

["The Physical Machines Page" on page 95](#)

["The Virtual Machines Page" on page 99](#)

["Business and Management Network Requirements" on page 28](#)

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

Replacing Physical Machines, Motherboards, NICs, or RAID Controllers

You replace motherboards, NICs, RAID controllers, and a physical machine (PM) in an everRun system while the system is running. You can remove 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 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 130.](#))

You remove a PM using the everRun PM **Remove** function, which deletes a PM from the everRun system's database. The everRun system then waits while an additional PM completes the process of joining the system.

If you replace a PM or a component, use vendor instructions, but first read ["Physical Machine Hardware Maintenance Restrictions" on page 270.](#)



Warning: This procedure deletes all software you may have installed on the PM and all PM configuration information you may have entered before the replacement procedure. After you complete this procedure, you must manually re-install all your software and reconfigure the PM to match your original settings.

Prerequisite: Obtain installation software for the everRun 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 -xzvf everRun_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 Software" on page 36](#).

Prerequisites: If you are replacing a PM, prepare the new PM:

1. Configure networks. See ["Network Architecture Overview" on page 19](#).
2. Configure storage. See ["Storage Requirements" on page 26](#).
3. Connect power. See ["Connecting Power" on page 35](#).
4. Configure the firmware (BIOS or UEFI). See ["Configuring Settings in the Firmware Setup Utility" on page 40](#).

To remove and 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 displays **running (in Maintenance)**, click **Shutdown** and then **OK**.
4. After the PM has shut down, click **Remove** () and respond appropriately to the confirmation

message. An alert message appears if removal conditions are not met.

If you confirm that you want to remove the PM, the everRun software deletes the PM from the everRun system, and displays a message that the system has successfully deleted the PM.

To replace the PM, continue with the steps below.

5. Manually power off the old PM.
6. 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.
7. 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.
8. Manually power on the PM. As the PM powers on, enter the firmware (BIOS or UEFI) setup utility, and set the Optical Drive as the first boot device.
9. Either mount the ISO image or insert the DVD into the PM.
10. 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 50](#) for reference. Although that topic refers to the "second PM," in this instance, it applies to the replacement PM.

11. If any disk contains previously installed data, the following message appears and the system reboots (if no disk contains previously installed data, installation continues with the next step):

Rebooting because disks XXX were erased.

When the reboot is complete, the boot menu reappears and you must continue with Step 10, above.
12. 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)**.
13. When prompted to configure **ibiz0**, select **Automatic configuration via DHCP** or **Manual Configuration (Static Address)**. (The installation software configures **priv0** automatically.)
14. When installation is complete, the PM ejects the install DVD (if used) and reboots.
15. 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 (in Maintenance)**.

16. Assign logical disks from the new PM to storage groups on the everRun system, as described in ["Assigning a Logical Disk to a Storage Group" on page 124.](#)

Notes:

- When the new PM joins the everRun system, the system automatically adds the secondary everRun system disk to the Initial Storage Group; however, the system does not assign any other logical disks from the new PM to existing storage groups.
- If you assigned logical disks to the Initial Storage Group or other storage groups on the first PM, you must manually add matching logical disks from the new PM to the same storage groups; otherwise, the everRun system cannot fully synchronize.

17. When you are ready to bring the replacement PM online, click **Finalize** to exit maintenance mode. Verify that both PMs return to the **running** state and that the PMs finish synchronizing. The initial synchronization may take minutes or hours depending on your configuration, including the amount of storage and the number of VMs.
18. Manually reinstall applications and any other host-level software and reconfigure the PM to match your original settings.

Related Topics

["Maintenance Mode" on page 125](#)

["Maintaining Physical Machines" on page 269](#)

["The everRun Availability Console" on page 58](#)

["Physical Machines and Virtual Machines" on page 9](#)

["The Physical Machines Page" on page 95](#)

Replacing a Failed Physical Machine Connected to External Storage

You can replace a failed physical machine (PM) that is connected to external storage while the system is running. You begin by unmapping and deleting all LUNs on the external storage array that are mapped to the failed PM. After the failed PM is replaced, you need to update the `lvm.conf` file, configure multipath, create new LUNs on the external storage systems, scan for new LUNs, configure `initrd` for multipath, and then place the LUNs in the proper storage group(s).

For detailed information, access the Knowledge Base to search for the article *Replacing a Failed Physical Machine Connected to External Storage* (KB-4298). See ["Accessing Knowledge Base Articles"](#) on page 481.

Related Topics

["External Storage"](#) on page 17

["Managing External Storage"](#) on page 279

["Maintaining Physical Machines"](#) on page 269

Upgrading Both Physical Machines In a Running System



Prerequisite: Before you upgrade to new physical machines, see ["Physical Machine Hardware Maintenance Restrictions"](#) on page 270.

To upgrade to new physical machines

1. Upgrade everRun 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 274.
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 19.

Related Topics

["Maintenance Mode"](#) on page 125

["Maintaining Physical Machines"](#) on page 269

["The everRun Availability Console"](#) on page 58

["Physical Machines and Virtual Machines"](#) on page 9

["The Physical Machines Page"](#) on page 95

9

Chapter 9: Managing External Storage

Manage external storage by performing these tasks:

- ["Adding or Removing External Storage LUNs" on page 279](#)
- ["Configuring External Storage in a Newly Installed Release 7.4.x System" on page 280](#)
- ["Configuring Linux Multipathing" on page 281](#)
- ["Removing External Storage Disks" on page 285](#)
- ["Updating initrd for Multipath" on page 285](#)
- ["Upgrading a System with External Storage from Release 7.3.4 to 7.4.x" on page 284](#)

For information on maintaining and troubleshooting physical machines with external storage, see the following:

- ["Recovering a Failed Physical Machine with External Storage " on page 132](#)
- ["Replacing a Failed Physical Machine Connected to External Storage" on page 277](#)

For general storage information, see:

- ["External Storage" on page 17.](#)
- ["everRun Storage Architecture" on page 15](#)

Adding or Removing External Storage LUNs

Use your storage array's software to map a new LUN to the everRun system or to remove a LUN. However, everRun software does not automatically detect the addition or removal of external storage LUNs. In order for everRun to recognize the changes, you must configure the host OS (CentOS):

- To add LUN0, you first remove LUNZ.
- To add LUN1 (or greater), you issue `lsscsi`, `echo`, `multipath`, and `dracut` commands.
- To remove LUNs, you issue `multipath`, `echo`, and `dracut` commands.

For detailed information, access the Knowledge Base to search for the article *Adding or Removing External Storage* (KB-4299). See ["Accessing Knowledge Base Articles" on page 481](#).

Related Topics

["External Storage" on page 17](#)

["Managing External Storage " on page 279](#)

["Replacing a Failed Physical Machine Connected to External Storage" on page 277](#)

Configuring External Storage in a Newly Installed Release 7.4.x System

After you install an everRun system, you can connect and configure external storage. everRun systems support Fibre Channel and iSCSI storage systems.

Prerequisites: Before performing the procedure below, you need to::



- Install the system hardware and software. For information on installing the hardware, see the manufacturer's instructions. For information on installing the everRun software, see ["Software Installation" on page 34](#).
- If the external storage is Fibre Channel, you need to install the FC HBAs in each physical machine (PM) and map each HBA WWN properly to the SAN switches (if used) and the external storage array.
- Physically connect the external storage according to the manufacturer's instructions.

To configure external storage in a newly installed Release 7.4.x system

Perform the following steps on node0, first, and then on node1:

1. Edit the `/etc/lvm/lvm.conf` file, adding the boot disk and data disk(s) to the filter entry. The following steps are an example:
 - a. Uncomment the `filter` entry.
 - b. Add the boot disk `a/sda/` and any internal data disk(s) installed in the system to the filter

entry. The following example is for a system with a data disk `sdb`:

```
filter = [ "a/sda/", "a/sdb/" ]
```

2. Configure multipath, as described in ["Configuring Linux Multipathing" on page 281](#).
3. Configure external storage, as described in ["Adding or Removing External Storage LUNs" on page 279](#) as well as the instructions provided by your external storage system vendor.

If the system is connected to iSCSI external storage, copy (or make note of) the new iSCSI `'iqn'` from the PM you replaced and update the PM's host entry in the external storage array with the new `'iqn'`. You must perform this step before you can map the new LUNs to the new PM

4. Create new LUNs on the external storage array and map them to new PM.
5. Create new storage groups.

For example, if you map two LUNs, one 250 GB and one 750 GB, to the node, you could create the new storage groups `StorageGroup_250GB` and `StorageGroup_750GB`.

6. Rescan the node in order to discover new LUNs by issuing `echo` commands (for additional information, see ["Adding or Removing External Storage LUNs" on page 279](#)).

The following commands are examples:

```
echo "-- --" > /sys/class/scsi_host/host1/scan
echo "-- --" > /sys/class/scsi_host/host6/scan
```

After the scan, you can view the new LUNs in the everRun Availability Console by navigating to the Storage tab of the **Physical Machines** page. See ["The Physical Machines Page" on page 95](#).

7. Place the LUNs into the appropriate storage group (created in Step 5, above). For example, place the 250 GB LUN in `StorageGroup_250GB`.

Repeat this procedure on the other node.

Related Topics

["Getting Started" on page 23](#)

["External Storage" on page 17](#)

["Managing External Storage" on page 279](#)

["Replacing a Failed Physical Machine Connected to External Storage" on page 277](#)

Configuring Linux Multipathing

You must configure the everRun system to use Linux multipathing if the system includes external storage. Multipathing enables the system to provide multiple paths to external storage systems. When everRun software is installed, the device-mapper-multipath package is also installed. However, everRun software does not configure multipath. Follow the instructions provided in the CentOS documentation to configure multipath. See https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/DM_Multipath/.

For everRun software to function properly, you must perform the required procedures below on both nodes after configuring multipath.

To set path_selector, blacklist drbd, and then set no_path_retry

1. Set up the initial `/etc/multipath.conf` file by issuing the following command:

```
mpathconf --enable --user_friendly_names y --find_multipaths y
--with_multipathd y
```

2. Edit the `/etc/multipath.conf` file, making the following changes:

- a. Add the entry `path_selector "round-robin 0"` to the `defaults` section, as in the following example.

```
## Use user friendly names, instead of using WWIDs as
names.

defaults {

    user_friendly_names yes

    find_multipaths yes

    path_selector "round-robin 0"

}
```

- b. Add or modify the `blacklist` entry (at the end of the `multipathd` file) to include `drbd`, so that the entry is identical to the following:

```
blacklist {

    devnode "^(ram|zram|raw|loop|fd|md|dm-
|sr|scd|st|drbd)[0-9]*"

}
```

- c. Add the following `devices` entry just below the `blacklist` line:

```
devices {
    device {
        vendor "DGC"
        product "VRAID"
        no_path_retry "0"
    }
}
```



Note: The example shown above is for an EMC storage array. The values for `vendor` and `product` vary according to the manufacturer and model of your storage array.

- d. Save the `/etc/multipath.conf` file.
3. To restart `multipathd` and to ensure that it can start at boot time and is functioning properly, issue the following commands in the order given:

```
systemctl restart multipathd
systemctl enable multipathd
systemctl status multipathd
```

4. To verify that `no_path_retry` is set to fail, issue the following command:

```
multipathd show config
```

The following sample lines appear at the end of output for an EMC VNX configuration:

```
device {
    vendor "DGC"
    product "VRAID"
    product_blacklist "LUNZ"
    path_grouping_policy "group_by_prio"
```

```
path_checker "emc_clariion"  
features "0"  
hardware_handler "1 emc"  
prio "emc"  
failback immediate  
rr_weight "uniform"  
no_path_retry "fail"  
retain_attached_hw_handler yes  
detect_prio yes  
}
```

Related Topics

["External Storage" on page 17](#)

["Managing External Storage " on page 279](#)

Upgrading a System with External Storage from Release 7.3.4 to 7.4.x



Caution: The **Qualify** procedure will not allow an upgrade from everRun Release 7.3.4 to 7.4.0 on a system with external storage. You must call your Stratus service representative to enable the upgrade.

Upgrading from everRun Release 7.3.4 to Release 7.4.x on a system with external storage requires additional steps before and after the software upgrade. For instance, before you begin the upgrade, you must first shutdown all VMs and delete all deployed VCDs. In addition, the upgrade includes updating `initrd` as well as modifying the `multipath.conf` and `lvm.conf` files for systems with Fibre Channel or iSCSI external storage. The upgrade for systems with iSCSI external storage also includes stopping and restarting the `everrun-sim` service, remapping LUNs, and using `iscsiadmin` for various tasks.

For the latest technical information, contact your Stratus service representative and refer to KB-4273.

Related Topics

["Managing External Storage " on page 279](#)

["Upgrading everRun Software" on page 115](#)

["External Storage" on page 17](#)

Removing External Storage Disks

To remove external storage disks, you need to unmap the LUN(s) from the external storage array, rescan the disks, delete the devices and the device nodes, issue `dracut`, and reboot the node, performing these tasks on `node0` and `node1`. Finally, repair the disk(s).

For detailed information, access the Knowledge Base to search for the article *Removing External Storage Disks* (KB-4335). See ["Accessing Knowledge Base Articles" on page 481](#).

Related Topics

["External Storage" on page 17](#)

["Managing External Storage " on page 279](#)

["The Physical Machines Page" on page 95](#)

Updating initrd for Multipath

For some external storage procedures, you need to update the initial RAM disk image (`initrd`) for multipath using `dracut`.

For detailed information, access the Knowledge Base to search for the article *Updating initrd for Multipath* (KB-4295). See ["Accessing Knowledge Base Articles" on page 481](#).

Related Topics

["Managing External Storage " on page 279](#)

["External Storage" on page 17](#)

Part 2: Supporting Documents

See the following support documents for release information, and reference and troubleshooting information.

- ["everRun Release 7.4.3.2 Release Notes" on page 288](#)
- ["everRun Command Line Interface Reference" on page 308](#)
- ["System Reference Information" on page 474](#)
- ["SNMP" on page 484](#)

10

Chapter 10: everRun Release 7.4.3.2 Release Notes

These Release Notes are for everRun release 7.4.3.2 (updated at 11:10 AM on 1/16/2019). See the following sections:

- [Accessing Stratus Knowledge Base Articles](#)
- [Bug Fixes](#)
- [New Features and Enhancements](#)
- [Important Considerations](#)
- [Known Issues](#)
- [Documentation Updates](#)
- [Getting Help](#)

Notes:



- For the latest technical information and updates, please refer to the English version of the *everRun User's Guide* at the **Downloads** page at <http://www.stratus.com/go/support/everrun> and the **Knowledge Base** in the **Stratus Customer Service Portal** at <https://support.stratus.com> (see [Accessing Stratus Knowledge Base Articles](#) for more information).
- For information about documentation changes that were made late in the release and may appear only on the StrataDOC web site, see [Documentation Updates](#).

Accessing Stratus Knowledge Base Articles

The **Stratus Customer Service Portal** provides a searchable **Knowledge Base** with technical articles about all Stratus products, including everRun. In some cases, the Release Notes directly reference these Knowledge Base articles (for example, KB-*nnnn*). You can access the Customer Service Portal and the Knowledge Base articles by using your existing service portal credentials, or by creating a new user account, as follows.

To access the Knowledge Base

1. Log on to the **Stratus Customer Service Portal** at <https://support.stratus.com>

If needed, create a new account as follows:

- a. Click **Register Account**.
- b. Enter your company email address and contact info, and click **Register**.

Your company email address must include a domain name (for example, stratus.com) for a company that is a registered customer of Stratus.

- c. Click the link in the email that you receive from Stratus.
- d. Enter a new password and finish configuring your account.

If you need assistance creating an account, contact your authorized Stratus service representative.

2. In the service portal, click **Knowledge Base** in the left pane.
3. In the **Keyword Search** box, enter keywords associated with the information you need, and then click **Search**.

To search for an article by its KB-*nnnn* number, click **Advanced Search**. Next to **Search by ID**, type the article ID number (*nnnn*) and click **Display**.

Bug Fixes

Fixed in everRun Release 7.4.3.2

- EV-47797: Implement AX fault handler enhancements and a related thread monitoring service (wdog) to recognize and handle thread hangs.
- EV-47943: Missing process count in /proc/stat error in /var/log/messages.
- EV-48153: Guest lost running netstall.sh with thread hangs invoked.
- EV-48178: Spine (System Manager) does not start the AX-QEMU after it was shut down due to CoreFT Fault Handler.

- EV-48180: Node1 QEMU is not restarted after it crashed.
- EV-48182: QEMU will not restart after the guest crashed.
- EV-48262: If a guest has a network down or disabled, thread hangs can take out the only AX that can keep the guest up and/or on the network.

Fixed in everRun Release 7.4.3.1

- EV-47564: Address a CPU hard lockup caused by KVM.
EV-47564 fixes a regression introduced by Red Hat in CentOS 7.2 (tracked by Red Hat under https://bugzilla.redhat.com/show_bug.cgi?id=1337667). This regression can cause guest delays and guest hangs.
- EV-47605: Revert a fault handling spit brain avoidance change.
EV-47605 removes a change in the timing of the Stratus fault handler code that was introduced in Release 7.4.0.

Fixed in everRun Release 7.4.3.0

- EV-44796: QEMU may crash on the active AX when clicking the Console button, or during shutdown.
- EV-45604: systemd-udev may fail to rename a network interface name on node startup.
- EV-45661: VM network may encounter a blackout after a migration if tagged VLANs are used.
- EV-45778: Kit upgrade Qualify script fails to check alternate boot disks for sufficient available space.
- EV-45779: Install aborts on system with Adaptec Series 8 12G SAS/PCIe RAID Adapter (vendor id: ASR8160).
- EV-45780: A spine restart may lead to unnecessary VM failover.
- EV-45781: After system shutdown and removing a network adapter, VM will not start due to non-redundant business network.
- EV-45784: VMs may temporarily run on a failing node.
- EV-45785: VM can fail to boot after system shutdown and powering up secondary node.
- EV-45787: Unable to configure valid volume container size when copying a VM.

- EV-45880: Disk failures on both nodes can result in unrecoverable faulted volumes even after the disks are repaired.
- EV-45881: Disk failure on one node can leave VM volumes in a failed state even after the disk is repaired.
- EV-45882: Install may hang when entering network information.
- EV-45883: Unable to start a stopped VM that has a good boot volume but one or more broken data volumes.

Fixed in everRun Release 7.4.2.0

- EV-2707: DVD replace operation allows user to replace a PM without first removing the node. This can lead to a corrupt cluster.
- EV-3019: Apache Tomcat version can be displayed.
- EV-3748: Unable to enable a guest data volume after a storage group fault.
- EV-44216: Address the Dirty COW (CVE-2016-5195) vulnerability.
- EV-44323: VM may fail to boot or boot but is degraded.
- EV-44354: A-link disconnects can incorrectly trigger "disconnected link condition" alerts for business networks.
- EV-44419: The nodes .htaccess file is visible by specifying systemip/.htaccess in a web browser.
- EV-44422: Invalid application consistent snapshots of Window 2008 32 bit and W2003 32 bit VMs may not be displayed as such in the eAC.
- EV-44450: Failure to start an AX will result in a degraded VM.
- EV-44494: Brief transient A-Link disconnects can be seen during node maintenance operations.
- EV-44562: A new Broadcom driver is needed for a new set of NetXtreme Broadcom Adapters.
- EV-44636: A W2K12 guest will lose about 4 seconds after each transition.
- EV-44678: Different mirror copy directions on two sides.
- EV-44679: Changing the Network Role on a shared network and rebooting can lead to a network loop bringing down the network.
- EV-44680: Swap warning alert may specify wrong node.
- EV-44681: Disk replication state may incorrectly remain inconsistent.

- EV-44684: Additional logging of the upgrade qualify processing is needed.
- EV-44705: Address openssl CVEs 2016-6303 and 2016-6306.
- EV-44715: Iptables and Ip6tables services are not available on everRun running 7.4 and later.
- EV-44738: Deselecting a data volume when taking a snapshot results in an invalid snapshot and an alert.
- EV-44744: A P2V'ed VM on everRun 7.4.1.x will fail to boot with a BSOD.
- EV-44762: VM guest lost as a result of QEMU crashes.
- EV-44771: User logins and logouts are not recorded in the audit
- EV-44774: Network failover after segmentation does not work as expected.
- EV-44809: AVCLI ad-join/ad-remove do not work.
- EV-44933: Business and ALinks may not be reported as degraded or broken and alerts generated when they should.
- EV-44999: VM running very slow on the active node if the secondary node has a power loss.
- EV-45010: Unable to create a VM from a snapshot due to timeout.
- EV-45017: The eAC may become unresponsive in everRun 7.4.1.1.
- EV-45044: Exclude the multicasts generated from everRun systems from network health statistics.
- EV-45082: Remove the kit upgrade requirement that existing partitions must be aligned on a 1 MB boundary.
- EV-45108: The guest VM may crash or network traffic may stop for over 20 seconds during deletion of a large snapshot, for example over 200GB.
- EV-45113: Audit log display in the eAC may be empty.
- EV-45114: A node may fail to join due to a time out. This can be seen during an upgrade, recover, or replace.
- EV-45120: The Import/Restore wizard will hang if there is not enough space in any storage group to create volumes.
- EV-45148: Create VM may fail when only one node is present.
- EV-45215: Update everRun host OS from CentOS 7.2 to CentOS 7.3.

- EV-45237: One or more VMs on a node fails to boot when starting multiple VMs at same time.
- EV-45314: On rare occasions the everRun install or upgrade may fail in the partition setup process when trying to configure volume groups.
- EV-45434: Update the HP kmod-hpsa driver to the latest CentOS 7.3 version (kmod-hpsa-3.4.18-105.rhel7u3.x86_64).

Fixed in everRun Release 7.4.1.2

- EV-44995: VM running very slow on the active node if the secondary node has a power loss.
- EV-45000: VM guest lost as a result of QEMU crashes.
- EV-45001: Network blackout after a migration or failover if there is time drift between nodes.
- EV-45002: A W2K12 guest may lose about 4 seconds after each transition.
- EV-45014: The eAC may become unresponsive in eE 7.4.1.1.

Fixed in everRun Release 7.4.1.1

- EV-44882: Add preservation of UI customization to the upgrade hook scripts.
- EV-44835: Business network disconnected Alert is incorrectly created after pulling A-Link cables.
- EV-44834: Transient A-Link disconnects may be seen in everRun Availability Console when putting nodes into and taking nodes out of maintenance.

Fixed in everRun Release 7.4.1.0

- Updated the everRun VirtIO drivers to a newer version. To ensure the proper operation of your VMs, update the VirtIO drivers in your Windows-based VMs to the latest version, as described in "[Updating the VirtIO Drivers \(Windows-based VMs\)](#)" on page 215.
- EV-44355: VM may crash if the VM is restarted while a sync is in progress.
- EV-44327: Switched boot fails with boot hang or guest moved used index error. Guest is degraded.
- EV-44161: Remove DNS requirement. The requirement for at least one DNS server has been lifted.
- EV-44328: Install vioserial driver on Windows 10 VM creation. Prior to this fix, the driver had to be installed manually after the VM was created.

- EV-44288: Occasional guest network dropouts when using the Windows VirtIO drivers supplied with everRun Release 7.4.0.x. The VirtIO drivers have been updated to a newer version that addresses this issue.

Fixed in everRun Release 7.4.0.1

- EV-3160: Enhance the everRun Release 7.4.0.0 pre-upgrade script to apply policy fixes to the 7.3.4 node that improve upgrade.
- EV-44119: Increase the value for maximum memory per VM from 213.3 GB to 256 GB.

New Features and Enhancements

New in everRun Release 7.4.3.0

You can obtain information about alerts, audit logs, nodes, VMs, and volumes by issuing the `snmptable` command. For information, see ["Obtaining System Information with snmptable" on page 484](#).

New in everRun Release 7.4.2.0

Platform and Host OS support includes these Intel® Xeon® processors (also listed in ["Physical Machine System Requirements" on page 476](#)):

- Scalable processors:
 - Gold 6XXX, 5XXX
 - Silver 4XXX
 - Bronze 3XXX
- E3-1XXXV6

An Ubuntu 14.04 AVCLI client has been added and can be downloaded from the **Downloads** page at <http://www.stratus.com/go/support/everrun>.

An everRun system provides a *network segmentation detection mechanism*. For information, see ["Network Segmentation Fault Detection and Remediation" on page 21](#).

everRun Release 7.4.2.0 and later releases support CentOS and Red Hat Releases 6.9 and 7.3 as well as the SLES 11 SP4 64-bit and SLES 12 SP2 64-bit versions of SUSE Linux Enterprise Server. See ["Compatible Guest Operating Systems" on page 474](#).

The kit upgrade requirement that existing partitions must be aligned on a 1 MB boundary has been removed.

New in everRun Release 7.4.1.0

everRun Release 7.4.1.0 and later releases support Microsoft Windows Server 2016 Essentials, Standard, and Datacenter (64-bit) guest operating systems for VMs, as listed in "[Compatible Guest Operating Systems](#)" on page 474.

New in everRun Release 7.4.0.0

- Platform and Host OS Support
 - Intel Broadwell processors (supported in Release 7.3.4.0 or higher)
 - Intel Xeon E3-1XXX v4
 - Intel Xeon E3-1XXX v5
 - Intel Xeon E5-1XXX v4
 - Intel Xeon E5-2XXX v4
 - CentOS 7.2 host operating system
- Performance Improvements
 - Page Modification Logging (PML) with Intel Broadwell processors
 - 4K sector disks
- Host OS Installation
 - Unified Extensible Firmware Interface (UEFI) support
 - Installation from USB device (for example, a thumb drive)
 - PXE host installation
 - Self-verifying checksum for installation ISO
- Host OS Upgrade
 - Upgrade supported from Release 7.3.4.0 to 7.4.0.0 or higher
 - Pre-upgrade check (**Qualify** button on **Upgrade Kits** page)

- Virtual machines
 - Support for new guests
 - Windows 10
 - RedHat/CentOS 6.7, 6.8, 7.1, and 7.2
 - Ubuntu 16.04
 - Support for 28 VMs (up to 8 of these can be FT VMs)
 - ISO image repository for VM creation
 - Console keyboard selection (English, German, Japanese)
 - Expand volume size in everRun Availability Console
 - Guest CD/DVD access
 - Detach/reattach boot volumes for troubleshooting
 - Copy a shutdown VM (without snapshot)
 - Export a shutdown VM (without snapshot)
- Network
 - Jumbo Frames on business networks
 - Enable/disable network interfaces (AVCLI)
- Storage
 - 4K sector size for logical disks, storage groups, and VM data volumes
 - Enable/disable storage devices (AVCLI)
- everRun Availability Console
 - Java eliminated from VM wizards
 - Grayed-out buttons for inactive features
 - Guest OS displayed in VM details
 - Disk sync progress/direction on a per VM basis displayed in VM details

- Audit Log Updates
 - Log VM events and new user creation
 - Internationalized logs
- e-Alerts
 - Alternate SMTP port
 - TLS encryption
 - Include asset ID in e-Alerts
- Host Security
 - HTTPS-only option
 - IPTables port security option
 - Active Directory improvements
 - Host OS inactivity automatic logout
- Snapshots
 - Individual volumes
 - Enable/disable snapshots
 - Create VM directly from snapshot
- P2V client
 - Windows 8.x and Windows Server 2012 support
 - HA guest targets
 - Support for more than four volumes
 - Storage group placement control on target system

Features No Longer Supported in everRun Release 7.4.0.0 or Higher

- One View Console
- Disaster Recovery and simplex everRun nodes
- Nagios[®] monitoring solutions

Important Considerations

Some Browsers Unable To Connect A VNC When Using https

If you are connected to the everRun Availability Console using an **https** URL in a Microsoft Internet Explorer™ or Mozilla® FireFox® browser, and you click **Console** after selecting a running VM from the **Virtual Machines** page, the message **Unable to connect VNC** may appear. To enable the VNC connection, click the https link to the VNC console page in the upper right-hand corner of the masthead, and continue with the appropriate procedure below:

- In Internet Explorer, the Security Alert wizard appears:
 - a. Click **Continue to this website (not recommended)**.
 - b. Click **OK**.
- In FireFox, the **Your connection is not secure** window appears:
 - a. Click **Advanced**. A message about an invalid security certificate appears.
 - b. Click **Add Exception**. The **Add Security Exception** dialog box appears with the console's location in **Location**.
 - c. Click **Confirm Security Exception**.

The VNC console appears.

Upgrading From everRun Release 7.4.0.x

You can upgrade from everRun Release 7.4.0.x to Release 7.4.3.x by following the instructions in "[Upgrading everRun Software](#)" on page 115.



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

Upgrading From everRun Release 7.3.4.0 or Previous Releases of everRun

You can upgrade from everRun Release 7.3.4.0 to Release 7.4.x by following the instructions in "[Upgrading everRun Software](#)" on page 115. If you are upgrading from an everRun release prior to Release 7.3.4.0, you must upgrade to Release 7.3.4.0 first, and then upgrade to Release 7.4.x.

The upgrade to everRun Release 7.4.x requires no VM downtime unless your system has external storage, as summarized in "[Upgrading a System with External Storage from Release 7.3.4 to 7.4.x](#)" on page 284.



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

When upgrading from everRun 7.3.4.0 or previous releases to everRun Release 7.4.x, consider the following:

- One View Console and Disaster Recovery are not supported.

If your everRun system is registered with the One View Console, you must unmanage the system before upgrading. If your VMs are protected by Disaster Recovery, you must unprotect the VMs, and then unmanage the everRun system before upgrading. For information on completing these tasks, see the One View and Disaster Recovery online help.

- Nagios monitoring solutions are not supported.

everRun Release 7.4.x no longer supports Nagios monitoring which was a technology preview in everRun 7.3. Nagios add-ons are automatically removed during the upgrade process.

- Upgrading everRun systems with external storage requires assistance from Stratus



Caution: The **Qualify** procedure will not allow an upgrade from everRun Release 7.3.4 to 7.4.x on a system with external storage. You must contact Stratus Customer Service to enable the upgrade.

Upgrading to Release 7.4.x on a system with external storage requires additional steps before and after the software upgrade, as summarized in "[Upgrading a System with External Storage from Release 7.3.4 to 7.4.x](#)" on page 284. For example, you will need to shut down all VMs and delete all deployed VCDs to begin the process. For the latest technical information, please contact your Stratus service representative and refer to KB-4273. Also see [Documentation Updates](#) for information about other new and updated external storage topics.



Note: When adding or upgrading external storage attached to your everRun system, you must configure the everRun system to use Linux multipathing to provide redundant paths to your external storage systems. See "[Configuring Linux Multipathing](#)" on page 281.

- everRun Release 7.4.x requires more disk space than previous releases

everRun Release 7.4.x requires more disk space than previous releases to accommodate changes for the host operating system. The amount of space required can be up to 10 GiB per logical system disk. To verify that your Release 7.3.4.0 system meets the upgrade requirements, upload the Release 7.4.x upgrade kit on the **Upgrade Kits** page of the everRun Availability Console. Uploading the kit gives you access to the **Qualify** button that you can click to run qualification checks. (Note that the **Qualify** button is available only in 7.3.4.0 or higher.) For more information about using the **Qualify** button and upgrading your system, see ["Upgrading everRun Software" on page 115](#).

- Windows-based VMs require Red Hat VirtIO driver updates after upgrade.

After upgrading to Release 7.4.x, update the VirtIO drivers in your Windows-based VMs to the latest supported versions to ensure the proper operation of your VMs. Download and install the supported drivers from the everRun support site, as described in ["Updating the VirtIO Drivers \(Windows-based VMs\)" on page 215](#).



Caution: RHEL 7.0 and CentOS 7.0 Virtual Machines must use kernel version 3.10.0-123.8.1 or later. If using an earlier kernel version there is a risk that the VM may hang.

- Any manually added RPMs or customized host OS configuration settings are lost.

The upgrade kit for Release 7.4.x performs the upgrade by reinstalling the everRun software on each node and restoring any settings associated with everRun features and management; however, the upgrade does not preserve any supplementary packages (for example, Dell OpenManage) or customized settings (for example, SNMP monitoring) that you may have implemented. If applicable, make note of any customizations and manually restore them after the upgrade.

- Any manually added IPtables port filtering changes are lost.

everRun Release 7.4.x includes improved IPtables port filtering, which you manage in the everRun Availability Console. As a result of the changes, upgrading will overwrite any settings you have manually entered in the host OS. If applicable, make note of any customizations and manually restore them after the upgrade. Use only the everRun Availability Console to configure IPtables filtering in Release 7.4.x or higher.

Using 4K Disks for Better Performance

everRun Release 7.4.0.0 or higher supports 4K sector size disks in native mode. Stratus recommends using 4K disks for better performance; however, observe the restrictions described in ["Storage Requirements" on page 26](#).

Do Not Update CentOS Host OS Directly From CentOS

Do not update the CentOS host OS of the everRun system directly from CentOS. Use only the CentOS release that is installed along with everRun software.

Optimizing Performance of A-Link Networks

Enabling jumbo frames on A-Link networks by changing their Ethernet frame MTU size from the default 1500 bytes to 9000 bytes may improve VM performance and reduce host processing overhead. See [KB-4262](#) for instructions.

Status of RAID Physical Disks Is Not Monitored

everRun 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 Considerations

For important considerations about everRun systems, see ["Important Physical Machine and Virtual Machine Considerations" on page 478](#).

Known Issues

Avoid Removing a VM While Only One Node is Running

To avoid downtime, always ensure that both nodes are in the **running** state and that neither node is in maintenance mode or in the process of synchronizing before removing a VM.

If you remove a VM while one node is shut down in maintenance mode, the operation may stall with the error, "There was a problem deleting *vm-name*." To correct the problem, you must restart the only running node.

Invalid VM Names Prevent Some VM Operations from Executing

The everRun Availability Console may improperly validate VM names to ensure that they meet the software requirements. If you are creating, reprovisioning, or renaming a VM with an unsuitable name, the operation may fail to execute.

To avoid any problems with VM operations, ensure that your VM names meet the following requirements:

- A VM name can start only with a word or a number, and the name cannot include the special characters (for example, #, %, or \$).
- A VM name cannot use hyphenated prefixes like Zombie- or migrating-.
- A VM name cannot be longer than 85 characters.

Windows 2008 (pre-R2) Guests May Crash

To avoid guest crashes, do not install Windows 2008 pre-R2 guests on everRun systems. For more information, see [KB-1912](#).

Alert Status For a Failed Volume May Not Be Displayed Correctly

An alert indicating that a volume has failed on a specific node may have a Normal status indicator (green check mark) on the **Alerts** page even though the volume remains in a failed condition. In such a case, the **Repair Storage** button will still be visible. Click it to move the volume to another logical disk.

everRun May Not Detect Reconnection of Management Network Cable

If you boot a PM while its management network (ibiz0) cable is disconnected and then you reconnect the cable after the PM is running, everRun may not detect that the cable has been reconnected. For more information, see [KB-4271](#).

Information About PCI Devices May Not Be Displayed

Depending on a system's hardware configuration, information about PCI devices (and the storage attached to them) may not be displayed in the everRun Availability Console.

State of External Storage Not Reflected in Node's State

A PM's (node) state is not affected by the state of non-system dependent volumes (any volumes other than root, swap and shared.fs). As a result, external storage can be in a failed state and the node's state will not reflect that. To see the state of external storage, view the state of the VMs and volumes that are on external storage.

Resizing Logical Disks on External Storage Is Not Supported

You cannot extend an external storage LUN and resize the everRun logical disk it contains to use the additional space. If more disk space is needed, you must create a new, larger LUN on your external storage

system and then map it to the everRun system as described in ["Adding or Removing External Storage LUNs" on page 279](#).

VMs May Not Boot If One Node is Removed From System

If you remove one PM from the system (by clicking **Work On** and then clicking **Remove**), do **not** place the remaining PM into maintenance and then bring it back into service by just using the **Finalize** button. Doing so will prevent all VMs from being able to boot. If you must place the remaining PM into maintenance mode, reboot it before you return it to service. See ["Rebooting a Physical Machine" on page 127](#) for details.

Removing Snapshots Temporarily Prevents Some VM Operations

When you remove a snapshot on an everRun system, the system must coalesce the snapshot by merging it with the next oldest snapshot. **While the system is coalescing snapshots:**

- A user cannot create a new snapshot in the everRun Availability Console. If you try, an error indicates that the system is busy.
- A user cannot start the VM associated with the snapshot(s) if the VM is currently stopped. The **Start** button is temporarily unavailable on the **Virtual Machines** page of the everRun Availability Console.
- A user cannot perform tasks that require the storage space occupied by the snapshot(s) until the coalescing operation is complete and the snapshot(s) are finally removed from the volume container. For example, this could prevent you from resizing a volume.

Avoid removing snapshots if you have an immediate need to perform any of these operations. After you remove a snapshot, wait at least 10-15 minutes before attempting any of these operations, or retry the operation if needed. You may need to wait much longer depending on the size of your volumes, the amount of VM activity, and the number of snapshots you remove.

For information about how to monitor coalescing operations that are underway, see [KB-4272](#).

Creation of Snapshots Results in Volume Conversion from RAW to QCOW2 Format

If you create and delete a snapshot of a volume that is currently in RAW format, the everRun software automatically converts the volume from RAW to QCOW2 format. The conversion is expected because snapshots require QCOW2 format, but consider that changing the volume format may have performance implications for certain VM loads. You cannot convert a volume from QCOW2 back to RAW format.

New features in everRun Release 7.4.0.0 or higher allow you to copy VMs and export VMs that are stopped without requiring a snapshot. If routine snapshots will not be required, consider these alternatives to retain RAW format (and avoid growing the size of volume containers).

Moving an everRun System to a Different Subnet

If you need to configure the network of an everRun system to work on a different subnet (for example when shipping to another location or when reconfiguring network subnets), see [KB-4264](#) for instructions.

Enabling Log Files for the Windows QEMU Guest Agent May Cause VM Timeouts When Taking Snapshots

When configuring the Windows QEMU guest agent, do not enable the option to save a log file during snapshots. If the QEMU guest agent attempts to create a log file during a snapshot, it may cause VSS timeouts that prevent the snapshot from completing.

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 that use the Broadcom BCM57712 Ethernet hardware chip

Do Not Use the `ifdown` Command

Do not issue the `ifdown` command from an everRun 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.

Documentation Updates

This section lists documentation changes made late in the release.

Updates in everRun Release 7.4.2.0 Documentation

- The `avcli "ad-join"` on page 329 and `"ad-remove"` on page 330 commands include the `--secure-mode` option.
- The `avcli "node-shutdown"` on page 394 command description includes examples.

Updates in everRun Release 7.4.1.0 Documentation

- ["Compatible Guest Operating Systems" on page 474](#) includes the Oracle Linux 6.8 and 7.3 Red Hat Compatible (RHC) Kernels.
- ["Configuring Linux Multipathing" on page 281](#) provides correct reference for multipath documentation.
- The Intel Xeon E3-1XXX v5 Processor is a supported CPU. See ["Physical Machine System Requirements" on page 476](#).

Updates in everRun Release 7.4.0.0 Documentation

- Updated firmware setup information to indicate that a system with UEFI firmware always boots from the original software boot disk in ["Configuring Settings in the Firmware Setup Utility" on page 40](#)
- Updated instructions for mapping your keyboard before or after installation in ["Mapping Your Keyboard" on page 48](#)
- Updated instructions for creating bootable USB installation media for the everRun software on Windows-based systems in ["Creating Bootable USB Media" on page 37](#)
- Updated everRun installation procedures to describe the behavior if any disk contains previously installed data in ["Installing Software on the First PM" on page 45](#) and ["Installing Software on the Second PM" on page 50](#)
- Updated everRun upgrade procedure to describe disk space requirements and the new **Qualify** button in , see ["Upgrading everRun Software" on page 115](#)
- New procedure for updating VirtIO drivers in Windows-based VMs, which has also been corrected to show installation of the **viostor** driver instead of the **vioscsi** driver in see ["Updating the VirtIO Drivers \(Windows-based VMs\)" on page 215](#)
- Updated storage requirements to discuss the newly supported 4K sector size in ["Storage Requirements" on page 26](#)
- Updated procedures to indicate that you must manually add matching logical disks to storage groups in a newly installed secondary or replacement node in:

- ["Converting an Avance Unit to an everRun 7.x System" on page 173](#)
- ["Converting an everRun MX System to an everRun 7.x System" on page 164](#)
- ["Replacing Physical Machines, Motherboards, NICs, or RAID Controllers" on page 274](#)
- New and updated external storage topics to describe procedures for everRun Release 7.4.0.0 or higher in:



Note: In some cases, the contents of these topics have been moved to the Stratus Knowledge Base. See the associated KB-*nnnn* articles for detailed information.

- ["Adding or Removing External Storage LUNs" on page 279](#) (Updated, was "Configuring External Storage") (Moved to KB-4299)
- ["Configuring External Storage in a Newly Installed Release 7.4.x System" on page 280](#) (New)
- ["Configuring Linux Multipathing" on page 281](#) (Updated)
- ["Removing External Storage Disks" on page 285](#) (New) (Moved to KB-4335)
- ["Updating initrd for Multipath" on page 285](#) (New) (Moved to KB-4295)
- ["Upgrading a System with External Storage from Release 7.3.4 to 7.4.x" on page 284](#) (New) (For the latest technical information, please contact your Stratus service representative and refer to KB-4273.)
- ["Recovering a Failed Physical Machine with External Storage " on page 132](#) (New) (Moved to KB-4300)
- ["Replacing a Failed Physical Machine Connected to External Storage" on page 277](#) (New) (Moved to KB-4298)
- Updated information to explain that if you change the **Static System IP** address of the everRun system, the automatically-generated MAC addresses of the VMs will also change in ["Configuring IP Settings" on page 71](#) and ["Assigning a Specific MAC Address to a Virtual Machine" on page 263](#)
- New `--detach-boot-volume` and `--attach-boot-volume` options for the `avcli vm-reprovision` command in ["vm-reprovision" on page 452](#)

- Updated system limits to indicate that the maximum memory per FT or HA VM is 256 GB in Release 7.4.0.1 or higher (it was 213.33 GB in previous releases due to management overhead) in "[Virtual Machine Recommendations and Limits](#)" on page 478

Getting Help

If you have a technical question about everRun software, you can find the latest technical information and online documentation at the **Downloads** page at <http://www.stratus.com/go/support/everrun>. You can also search the **Knowledge Base** in the **Stratus Customer Service Portal** at <https://support.stratus.com>.

If you are unable to resolve your questions with these online resources, and the system is covered by a service agreement, please contact everRun Customer Support or your authorized Stratus service representative. For information, see the **everRun Support** page at <http://www.stratus.com/services-support/customer-support>.

11

Chapter 11: everRun Command Line Interface Reference

You can use the everRun 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 308](#)
- ["AVCLI Command Descriptions" on page 320](#)

AVCLI Command Overview

You can use the everRun 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 309](#)
- ["Installing the Linux Client" on page 309](#)
- ["Installing the Windows Client" on page 310](#)

The following topics explain how to use the AVCLI command interface:

- ["Using AVCLI" on page 311](#)
- ["Executing a Command" on page 311](#)
- ["Using AVCLI Help" on page 312](#)

The following topics are helpful for programmers using the AVCLI command interface:

- ["AVCLI Error Status" on page 313](#)
- ["XML Encapsulated Errors" on page 313](#)

- ["Error Checking" on page 314](#)
- ["Asynchronous Command Delay" on page 314](#)
- ["Output Formatting" on page 315](#)
- ["AVCLI Exceptions" on page 319](#)

Related Topics

["AVCLI Command Descriptions" on page 320](#)

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 **Downloads** page at <http://www.stratus.com/go/support/everrun>.
 - b. On the **Downloads** page, click **everRun** (if it is not already displayed) and then select the appropriate version.

- c. Scroll down to **Drivers and Tools** and then continue scrolling to **everRun Command Line Interface (AVCLI)**.
 - d. Select the **RHEL (64-bit) avcli client** and 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
conflicts with file from package avcli-1.0-0 file
/usr/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 **Downloads** page at <http://www.stratus.com/go/support/everrun>.
 - b. On the **Downloads** page, click **everRun** (if it is not already displayed) and then select the appropriate version.
 - c. Scroll down to **Drivers and Tools** and then continue scrolling to **everRun Command Line Interface (AVCLI)**.
 - d. Click **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 311](#).

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 312](#).

Executing a Command

Commands must include the DNS name or IPv4 address of the everRun 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 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 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:014
-> 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.throwNo
nExistentResource (ResourceDisambiguateServiceProvider.java:56)

at
com.avance.yak.cli.ResourceDisambiguateServiceProvider.getReso
urceId (ResourceDisambiguateServiceProvider.java:81)

at
com.avance.yak.cli.Command.findResourceId (Command.java:80)

```

```
at
com.avance.yak.cli.CommandWithUnparsedAmbiguousResourcesInvoke
Each.execute
(CommandWithUnparsedAmbiguousResourcesInvokeEach.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.ResourceDisambiguateServiceProvider.throwNo
nExistentResource (ResourceDisambiguateServiceProvider.java:56)

at
com.avance.yak.cli.ResourceDisambiguateServiceProvider.getReso
urceId (ResourceDisambiguateServiceProvider.java:81)

at
com.avance.yak.cli.Command.findResourceId (Command.java:80)

at
```

```
com.avance.yak.cli.CommandWithUnparsedAmbiguousResourcesInvoke
Each.execute
(CommandWithUnparsedAmbiguousResourcesInvokeEach.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 311](#).

Help

["help" on page 355](#)

Basic System Information

["audit-export" on page 333](#)

["audit-info" on page 334](#)

["unit-change-ip" on page 417](#)

["unit-configure" on page 419](#)

["unit-eula-accept" on page 420](#)

["unit-eula-reset" on page 421](#)

["unit-info" on page 422](#)

["unit-shutdown" on page 423](#)

["unit-shutdown-cancel" on page 424](#)

["unit-shutdown-state" on page 425](#)

["unit-synced" on page 426](#)

System Configuration

["callhome-disable" on page 335](#)

["callhome-enable" on page 336](#)

["callhome-info" on page 337](#)

["datetime-config" on page 338](#)

["dialin-disable" on page 346](#)

["dialin-enable" on page 347](#)

["dialin-info" on page 348](#)

["ealert-config" on page 350](#)

["ealert-disable" on page 352](#)

["ealert-enable" on page 353](#)

["ealert-info" on page 354](#)

["license-info" on page 362](#)

["license-install" on page 363](#)

["ntp-config" on page 397](#)

["ntp-disable" on page 398](#)

["proxy-config" on page 404](#)

["proxy-disable" on page 405](#)

["proxy-enable" on page 406](#)

["proxy-info" on page 407](#)

["snmp-config" on page 408](#)

["snmp-disable" on page 410](#)

["snmp-info" on page 411](#)

["timezone-config" on page 414](#)

["timezone-info" on page 415](#)

System User Management

["ad-disable" on page 326](#)

["ad-enable" on page 327](#)

["ad-info" on page 328](#)

["ad-join" on page 329](#)

["ad-remove" on page 330](#)

["local-group-add" on page 364](#)

["local-group-delete" on page 365](#)

["local-group-edit" on page 366](#)

["local-group-info" on page 367](#)

["local-user-add" on page 368](#)

["local-user-delete" on page 370](#)

["local-user-edit" on page 371](#)

["local-user-info" on page 373](#)

["owner-config" on page 401](#)

["owner-info" on page 402](#)

Managing Physical Machines

["node-add" on page 386](#)

["node-cancel" on page 387](#)

["node-delete" on page 389](#)

["node-info" on page 391](#)

["node-reboot" on page 392](#)

["node-recover" on page 393](#)

["node-shutdown" on page 394](#)

["node-workoff" on page 395](#)

["node-workon" on page 396](#)

["pm-clear-mtbf" on page 403](#)

Managing Alerts

["alert-delete" on page 331](#)

["alert-info" on page 332](#)

Diagnostic Files

["diagnostic-create" on page 341](#)

["diagnostic-delete" on page 342](#)

["diagnostic-extract" on page 343](#)

["diagnostic-fetch" on page 344](#)

["diagnostic-info" on page 345](#)

Kit Information

["kit-delete" on page 360](#)

["kit-info" on page 361](#)

Network/Storage Information

["disk-move-to-group" on page 349](#)

["image-container-info" on page 356](#)

["image-container-resize" on page 359](#)

["network-change-mtu" on page 382](#)

["network-change-role" on page 383](#)

["network-info" on page 384](#)

["node-config-prp" on page 388](#)

["node-delete-prp" on page 390](#)

["storage-group-info" on page 412](#)

["storage-info" on page 413](#)

["volume-info" on page 471](#)

["volume-resize" on page 472](#)

Creating Virtual CD/DVDs

["media-create" on page 375](#)

["media-delete" on page 376](#)

["media-eject" on page 377](#)

["media-import" on page 378](#)

["media-info" on page 380](#)

["media-insert" on page 381](#)

Managing Virtual Machines

["localvm-clear-mtbf" on page 374](#)

["ova-info" on page 399](#)

["ovf-info" on page 400](#)

["vm-ax-disable" on page 427](#)

["vm-ax-enable" on page 428](#)

["vm-boot-attributes" on page 429](#)

["vm-cd-boot" on page 430](#)

["vm-copy" on page 431](#)

["vm-create" on page 435](#)

["vm-create-from-snapshot" on page 440](#)

["vm-delete" on page 442](#)

["vm-export" on page 443](#)

["vm-import" on page 444](#)

["vm-info" on page 447](#)

["vm-network-disable" on page 448](#)

["vm-network-enable" on page 449](#)

["vm-poweroff" on page 450](#)

["vm-poweron" on page 451](#)

["vm-reprovision" on page 452](#)

["vm-restore" on page 456](#)

["vm-shutdown" on page 459](#)

["vm-snapshot-create" on page 460](#)

["vm-snapshot-create-disable" on page 462](#)

["vm-snapshot-create-enable" on page 463](#)

["vm-snapshot-delete" on page 464](#)

["vm-snapshot-export" on page 465](#)

["vm-snapshot-info" on page 467](#)

["vm-unlock" on page 468](#)

["vm-volume-disable" on page 469](#)

["vm-volume-enable" on page 470](#)

Related Topics

["AVCLI Command Overview" on page 308](#)

ad-disable

Usage

```
avcli ad-disable
```

Description

The `ad-disable` command disables Active Directory support.

ad-enable

Usage

```
avcli ad-enable
```

Description

The `ad-enable` command enables Active Directory support.

ad-info

Usage

```
avcli ad-info
```

Description

The `ad-info` command displays information about the Active Directory.

ad-join

Usage

```
avcli ad-join --username name [--password password] [--secure-mode true|false] domain
```

Description

The `ad-join` command joins the everRun system to the specified Active Directory domain and enables Active Directory support.

Options

<code>--username <i>name</i></code>	The user with rights to join to the specified domain.
<code>--password <i>password</i></code>	The password of the user with rights to join to the specified domain. If you do not give the password, you are automatically prompted for it.
<code>--secure-mode true false</code>	Enable (<code>true</code> , the default) or disable (<code>false</code>) secure mode.
<i>domain</i>	The name of the Active Directory domain to join.

Examples

```
$ avcli ad-join --username domain\administrator --password  
secret domain
```

```
$ avcli ad-join --username domain\administrator domain
```

ad-remove**Usage**

```
avcli ad-remove --username name [--password password] [--
secure-mode true|false] domain
```

Description

The `ad-remove` command removes the everRun system from the specified Active Directory domain and disables Active Directory support.

Options

<code>--username <i>name</i></code>	The user with rights to remove the everRun system from the specified domain.
<code>--password <i>password</i></code>	The password of the user with rights to remove the everRun system from the specified domain. If you do not give the password, you are automatically prompted for it.
<code>--secure-mode true false</code>	Enable (<code>true</code> , the default) or disable (<code>false</code>) secure mode.
<code><i>domain</i></code>	The name of the Active Directory domain from which the everRun system is to be removed.

Examples

```
$ avcli ad-remove --username domain\administrator --password
secret domain
```

```
$ avcli ad-remove --username domain\administrator domain
```

alert-delete

Usage

```
avcli alert-delete [alerts... | purge]
```

Description

The `alert-delete` command deletes specific alerts or optionally, all alerts.

Options

<code>alerts</code>	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>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.

disk-move-to-group

Usage

```
avcli disk-move-to-group disk... storage-group
```

Description

The `disk-move-to-group` command moves one or more logical disks to a storage group.

Options

<i>disk</i>	One or more disks to be moved.
<i>storage-group</i>	A storage group.

ealert-config**Usage**

```
avcli ealert-config [--ssl] [--tls] [--port port] [--sender
sender] [--username username] [--password password] --host
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. You can specify this option only if <code>--tls</code> is not specified.
<code>--tls</code>	Use TLS when communicating with the SMTP server(s). You can specify this option only if <code>--ssl</code> is not specified.
<code>--port <i>port</i></code>	Use this port number to connect to the SMTP server.
<code>--sender <i>sender</i></code>	Send email from this sender (name or IP address).
<code>--username <i>username</i></code>	Use this name for authentication on the specified SMTP host.
<code>--password <i>password</i></code>	Use this password for authentication on the specified SMTP host.
<code>--host <i>host</i></code>	The DNS or IP address of the SMTP server.
<i>recipients</i>	The list of email addresses to receive eAlert emails; required only when e-Alerts are enabled.

Examples

The following command configures email alerts to use the SMTP server `mail.my-domain.com` and to send alerts to the recipient `admin@my-domain.com`:

```
$ avcli ealert-config --host mail.my-domain.com admin@my-domain.com
```

The following command configures email alerts to use the SMTP server `mail.my-domain.com` and to send alerts the recipients `admin@my-domain.com` and `bob@my-domain.com`:

```
$ avcli ealert-config --host mail.my-domain.com admin@my-domain.com bob@my-domain.com
```

The following command configures email alerts to use the protocol SSL and the SMTP server `mail.my-domain.com`, and to send alerts to the recipient `bob@my-domain.com`, using the username `admin` and the password `secret` for authentication:

```
$ avcli ealert-config --ssl --username admin --password secret --host mail.my-domain.com bob@my-domain.com
```

The following command configures email alerts to use the protocol SSL and the SMTP server `mail.my-domain.com`, and to send alerts to the recipient `bob@my-domain.com`, using the username `admin` and no password for authentication:

```
$ avcli ealert-config --ssl --username admin --host mail.my-domain.com bob@my-domain.com
```

The following command configures email alerts to send alerts from the sender `sample@gmail.com`, using the protocol TLS on port 587 and the SMTP server `mail.my-domain.com`, and to send alerts to the recipient `bob@my-domain.com`, using the username `admin` and no password for authentication:

```
$ avcli ealert-config --tls --port 587 --sender sample@gmail.com --username admin --password secret --host mail.my-domain.com 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
```

image-container-info

Usage

```
image-container-info [image-container]
```

Description

The `image-container-info` command displays information about all image containers (also known as *volume containers*) or optionally, about only the specified image containers. Specifically, the command displays information about the portion of the image container that is available to the guest operating system.

Options

<i>image-container</i>	The name of the image container. If you do not provide this argument, the command displays information about all image containers.
------------------------	--

Examples

```
$ avcli image-container-info
image-container:
-> name : root
-> id : imagecontainer:o58
-> hasFileSystem : false
-> isLocal : true
-> size : 21,479,030,784
-> size-used : 21,479,030,784
-> storage-group : none
image-container:
-> name : root
-> id : imagecontainer:o31
-> hasFileSystem : false
```

```
-> isLocal : true
-> size : 21,479,030,784
-> size-used : 21,479,030,784
-> storage-group : none
image-container:
-> name : swap
-> id : imagecontainer:o36
-> hasFileSystem : false
-> isLocal : true
-> size : 2,151,677,952
-> size-used : 2,151,677,952
-> storage-group : none
image-container:
-> name : swap
-> id : imagecontainer:o66
-> hasFileSystem : false
-> isLocal : true
-> size : 2,151,677,952
-> size-used : 2,151,677,952
-> storage-group : none
image-container:
-> name : shared.fs_image_container
-> id : imagecontainer:o77
-> hasFileSystem : false
-> isLocal : false
```

```
-> size : 1,073,741,824
-> size-used : 1,073,741,824
-> storage-group : none
image-container:
-> name : win7_ent_x86_32_sp1
-> id : imagecontainer:o1360
-> hasFileSystem : false
-> isLocal : false
-> size : 2,684,354,560
-> size-used : 2,684,354,560
storage-group:
-> name : Initial Storage Group
-> id : storagegroup:o21
image-container:
-> name : boot-chom1
-> id : imagecontainer:o1690
-> hasFileSystem : true
-> isLocal : false
-> size : 42,949,672,960
-> size-used : 37,787,627,192
storage-group:
-> name : Initial Storage Group
-> id : storagegroup:o21
```

image-container-resize

Usage

```
image-container-resize --new-size size image-container
```

Description

The `image-container-resize` command increases the size of the image container; specifically, the portion that is available to the guest operating system. (An *image container*, also known as a *volume container*, is a system-wide container that holds volumes and snapshots.) You might want to increase the image container's size if you need to take snapshots and the container lacks sufficient free space to do so.

Options

<code>--new-size <i>size</i></code>	The new image-container size. By default, <i>size</i> is in megabytes, but you can specify standard qualifiers (for example, KB, K, MB, M, GB, or G).
<code><i>image-container</i></code>	The name of the image container.

Examples

```
$ avcli image-container-resize --new-size 40G boot-chom1
```

kit-delete**Usage**

```
avcli kit-delete kit_id
```

Description

The `kit-delete` command deletes the specified kit(s).

Options

<i>kit_id</i>	One or more upgrade kits to be deleted. The value is the kit ID.
---------------	--

For information on obtaining the value for *kit_id*, see the "[kit-info](#)" on page 361 command description.

Example

```
kit-delete kit:o24
```

kit-info

Usage

```
avcli kit-info [kit_id...]
```

Description

The `kit-info` command displays information about all kits (the default) or only the specified kits.

Options

<i>kit_id</i>	One or more upgrade kits about which to display information. The value is the kit ID.
---------------	---

Examples

You can issue the `kit-info` command to obtain the *kit-id* value for the ["kit-delete" on page 360](#) command. In the command output, the `id` field displays the *kit-id* value. In the following sample output of the `kit-info` command, the `id` field displays the value `kit:o24`:

```
[root@node0 ~]# avcli kit-info
-> name : unspecified
-> id : kit:o24
-> description : unspecified
-> version : 7.5.0-127
-> locked : false
```

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. Only users with the role **Administrator** (the group `admin`) can execute this command. Users with the role **Platform Admin** (the group `platform_admin`) or **Read Only** (the group `read-only`) cannot execute this command.

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`). Only users with the role **Administrator** (the group `admin`) can execute this command. Users with the role **Platform Admin** (the group `platform_admin`) or **Read Only** (the group `read-only`) cannot execute this command.

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`). Only users with the role **Administrator** (the group `admin`) can execute this command. Users with the role **Platform Admin** (the group `platform_admin`) or **Read Only** (the group `read-only`) cannot execute this command.

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. Only users with the role **Administrator** (the group `admin`) can execute this command. Users with the role **Platform Admin** (the group `platform_admin`) or **Read Only** (the group `read-only`) cannot execute this command.

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. Only users with the role **Administrator** (the group `admin`) can execute this command. Users with the role **Platform Admin** (the group `platform_admin`) or **Read Only** (the group `read-only`) cannot execute this command.

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. Only users with the role **Administrator** (the group `admin`) can execute this command. Users with the role **Platform Admin** (the group `platform_admin`) or **Read Only** (the group `read-only`) cannot execute this command.

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.

<i>group-name-or-sid</i>	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.
-------------	---

localvm-clear-mtbf**Usage**

```
avcli localvm-clear-mtbf
```

Description

The `localvm-clear-mtbf` command brings half of a VM back into service after it has been removed from service for failing too many times.

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://hostname/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.
--------------	---

media-insert

Usage

```
avcli media-insert --iso [--cdrom] [vm...]
```

Description

The `media-insert` command allows you to insert media into the specified virtual machines.



Caution: When you insert a VCD into a running, fault-tolerant (FT) VM, it prevents the everRun software from migrating the VM to a different physical machine in the event of a failure. To restore fault-tolerant operation, unmount and eject the VCD as soon as you finish using it.

Options

<code>--iso <i>name</i></code>	The ISO image to insert.
<code>--cdrom <i>name</i></code>	The CD-ROM device to insert. This value is optional if the VM has only a single CD-ROM device.
<code><i>vm</i></code>	The name of the VM into which the media is to be inserted.

network-change-mtu

Usage

```
avcli network-change-mtu name size
```

Description

The `network-change-mtu` command changes the MTU size of the specified network (an A-Link or business network, including the biz0 network) on everRun systems.

Options

<i>name</i>	The name of the network
<i>size</i>	The MTU size. Valid values are 1500 - 9000.

1. Shut down the VM and issue `network-change-mtu`.

To change the MTU of an A-Link, perform the following steps:

1. Issue `network-change-mtu`.
2. Cycle a node through maintenance mode, which enables the node to acquire the new MTU (the AX TCP sessions disconnect and reestablish after both nodes are back online).
3. Repeat the step above on the second node.

Examples

The following commands change the MTU size on the A-Link `priv0`.

```
$ avcli network-change-mtu priv0 4000
```

```
$ avcli network-change-mtu priv0 9000
```

The following commands change the MTU size on the business network `network0`, sometimes known as `biz0`.

```
$ avcli network-change-mtu network0 4000
```

```
$ avcli network-change-mtu network0 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-config-prp

Usage

```
avcli node-config-prp --nic1 adapter --nic2 adapter node
```

Description

The `node-config-prp` command configures a PRP adapter on the specified PM with two physical adapters.

You must run this command twice: once to configure the adapter on the first PM, and once to configure the adapter on the second PM.

Options

<code>--nic1 <i>adapter</i></code>	The name of a physical adapter.
<code>--nic2 <i>adapter</i></code>	The name of a physical adapter.
<code><i>node</i></code>	The PM containing the PRP adapter to be configured.

Examples

```
$ avcli node-config-prp --nic1 eth0 --nic2 eth1 node0
```

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-delete-prp

Usage

```
avcli node-delete-prp --name adapter node
```

Description

The `node-delete-prp` command deletes a PRP adapter on the specified PM.

You must run this command twice: once to delete the adapter on the first PM, and once to delete the adapter on the second PM.

Options

<code>--name <i>adapter</i></code>	The name of the adapter to delete.
<code><i>node</i></code>	The PM containing the adapter to delete.

Examples

```
$ avcli node-delete-prp --name ad0 node0
```

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-reboot**Usage**

```
avcli node-reboot [--wait] pm
```

Description

The `node-reboot` command reboots the specified PM.

Options

<code>--wait</code> <code>-w</code>	Wait for the command to complete.
<code><i>pm</i></code>	The PM to reboot.

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 [--force] [--wait] pm
```

Description

The `node-shutdown` command shuts down the specified PM. Before issuing `node-shutdown`, you must first place the node in maintenance mode. To do so, you can issue "[node-workon](#)" on [page 396](#) or use the everRun Availability Console.

Options

<code>--force</code> <code>-f</code>	Override the shutdown warning.
<code>--wait</code> <code>-w</code>	Wait for the command to complete.
<code><i>pm</i></code>	The PM (for example, <code>node1</code>) to shut down.

Examples

```
$ avcli node-workon node0  
$ avcli node-shutdown --force node0
```

node-workoff

Usage

```
avcli node-workoff [--wait] pm
```

Description

The `node-workoff` command takes the specified PM out of maintenance mode.

Options

<code>--wait</code> <code>-w</code>	Wait for the command to complete.
<code><i>pm</i></code>	The PM to remove from maintenance mode.

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.
-----------	--

Examples

```
$ avcli node-workon node0
```

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.

pm-clear-mtbf

Usage

```
avcli pm-clear-mtbf
```

Description

The `pm-clear-mtbf` command clears a PM's MTBF from the user interface.

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 recipient ...] [--
recipients-v1 recipient-v1 ...]
```

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 recipient ...</code>	The list of hosts to which to send traps using SNMP version 2c; required only when traps are enabled and no <code>--recipients-v1</code> are listed.
<code>--recipients-v1 recipient-v1 ...</code>	The list of hosts to which to send traps using SNMP version 1; required only when traps are enabled and no <code>--recipients</code> are listed.

Examples

The following example enables SNMP requests, and then traps and sends them to `host1` and `host2` using SNMP version 2c, and to `snmp.my-domain.com` and `snmp2.my-domain.com` using SNMP version 1.

```
$ avcli snmp-config --enable-requests --enable-traps --  
recipients host1 host2 --recipients-v1 snmp.my-domain.com  
snmp2.my-domain.com
```

The following example disables SNMP requests, enables traps, and sends them to localhost using SNMP version 2c.

```
$ avcli snmp-config --enable-traps --community public --  
recipients localhost
```

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-avoid-bad-node

Usage

```
avcli unit-avoid-bad-node true|false|reset
```

Description

By default, VMs automatically return to a node that has returned to service after the node has recovered from a recent failure or was in maintenance mode. In some situations, you may want to verify that the node is healthy before VMs return to it. To prevent VMs from automatically returning to such nodes, set the migration policy. To do so, use the `unit-avoid-bad-node` command or see ["Configuring the Migration Policy" on page 79](#).

The `unit-avoid-bad-node` command enables or disables the ability of VMs to automatically return to a node that has recently failed or been in maintenance mode. If the node is healthy, issue `unit-avoid-bad-node reset` to re-enable VMs to automatically return to it.

If you issue this command with no option, the command checks if the setting is enabled or disabled, and displays `yes` or `no` values for `Feature enabled`, `Keeping VMs on last good node`, and `Awaiting reset signal`. The following output is an example:

```
Avoid automatically moving VMs back to a node that recovered
after a failure:
```

```
-> Feature enabled : yes
```

```
-> Keeping VMs on last good node : yes
```

```
-> Awaiting reset signal : yes
```

Options

<code>true</code>	Enables VMs to automatically return to a node that has returned to service. .
<code>false</code>	Disables VMs from automatically returning to a node that has returned to service.
<code>reset</code>	Enables VMs held on the last good node to return to a healthy node recently returned to service.

unit-change-ip

Usage

```
avcli unit-change-ip --cluster-address IP_address [--static]
[--prefix prefix] [--node0-address IP_address] [--node0-
gateway IP_address] [--node1-address IP_address] [--node1-
gateway IP_address] [--dns-servers server_address ...]
```

Description

The `unit-change-ip` command changes the IP configuration of the management network for the everRun system specified by `--cluster-address IP_address`.

Options

<code>--cluster-address <i>IP_address</i></code>	The IP address of the everRun system.
<code>--static</code>	Explicitly sets the values, if specified, for <code>--prefix</code> , <code>--node0-address</code> , <code>--node0-gateway</code> , <code>--node1-address</code> , <code>--node1-gateway</code> , and <code>--dns-servers</code> ; otherwise, DHCP sets these values (except <code>--cluster-address</code>).
<code>--prefix <i>prefix</i></code>	The size of the network prefix. Values, in bits, are 8 (Class A), 16 (Class B), and 24 (Class-C).
<code>--node0-address <i>IP_address</i></code>	The IP address of node0.
<code>--node0-gateway <i>IP_address</i></code>	The IP address of the node0 gateway.
<code>--node1-address <i>IP_address</i></code>	The IP address of node1.

<code>--node1-gateway <i>IP_address</i></code>	The IP address of the node1 gateway.
<code>[--dns-servers <i>server_address ...</i>]</code>	One or two DNS server(s). The first IP address is for the primary DNS server. The second (optional) IP address is for the secondary DNS server.

Examples

```
avcli unit-change-ip --cluster-address 10.92.179.54
avcli unit-change-ip --cluster-address 10.92.179.54 --static -
-prefix 16 --node0-address 10.92.179.154 --node0-gateway
10.92.0.1 --node1-address 10.92.179.156 --node1-gateway
10.92.0.1 --dns-servers 134.111.24.250 134.111.24.251
```

unit-configure

Usage

```
avcli unit-configure
```

Description

The `unit-configure` command configures the everRun system. The command implements the initial configuration of an everRun system, as executed by the **Portal Restart Required** window that appears after entering network information when logging on to the everRun Availability Console for the first time (see ["Logging on to the everRun Availability Console for the First Time" on page 54](#)). The `unit-configure` command causes all physical machines to exit Maintenance mode.

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> <code>-w</code>	Wait for the command to complete.
--	-----------------------------------

vm-ax-disable

Usage

```
avcli vm-ax-disable --name name --nodes nodes
```

Description

The `vm-ax-disable` command disables the VM's instance on selected PMs.

Options

<code>--name <i>name</i></code>	The name or ID of a VM.
<code>--nodes <i>nodes</i></code>	The name(s) or ID(s) of PM(s) to disable.

Examples

Disable instance on node1 for a VM named MyVM.

```
$ avcli vm-ax-disable --name MyVM --nodes node1
```

vm-ax-enable**Usage**

```
avcli vm-ax-enable --name name --node node
```

Description

The `vm-ax-enable` command enables the VM's instance on selected PM.

Options

<code>--name <i>name</i></code>	The name or ID of a VM.
<code>--node <i>node</i></code>	The name or ID of PM to enable.

Examples

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-copy

Usage

```
avcli vm-copy --source-vm source --name name [--description
description] [--cpu number] [--memory memory] [--availability
level] [--copy-volumes volumes] [--add-volumes volumes] [--
keep-volumes volumes] [--interfaces networks] [--storage-group
group] [--vnc-keyboard-layout keyboard] [--no-auto-start]
```

Description

The `vm-copy` command copies a VM from the specified VM. If any parameter is not specified, the corresponding value from the source VM is used.

Options

<code>--source-vm <i>source</i></code>	The name or ID of the source VM.
<code>--name <i>name</i></code>	The name of the VM to create.
<code>--description <i>description</i></code>	The description for the new VM.
<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>--availability <i>level</i></code>	The availability level, high availability (<code>ha</code>) or fault tolerant (<code>ft</code>).
<code>--copy-volumes <i>volumes</i></code>	<p>The list of volumes to copy to the new VM. Specify volumes by configuration name or ID, boot volume first.</p> <p>To copy all of the volumes from the source VM to the new VM with the default values, keep this parameter as blank.</p> <p>A <i>volume</i> is made up of five components, separated by commas:</p>

	<ul style="list-style-type: none"> • Volume name or ID of the source volume; required. • Storage-group name or ID from which to carve storage. • Volume name of the new volume. • Volume disk image format (raw or qcow2). • Sector size of the volume (512 or 4096). <ul style="list-style-type: none"> ▪ The sector size is in bytes (B), 512 B by default. ▪ If the sector size of the storage group is 512 B, the sector size of the volume must be 512 B. ▪ If the sector size of the storage group is in 4096 B (4 kB), both 512 B and 4096 B are supported as the sector size of the volume. ▪ The boot volume must be presented with 512 B as the sector size.
<pre>--add-volumes <i>volumes</i></pre>	<p>The list of volumes to attach to this VM. A <i>volume</i> is made up of five components, separated by commas:</p> <ul style="list-style-type: none"> • Size of the volume; required. By default, the volume size is specified in megabytes, but you can use standard qualifiers such as KB, MB, GB, and TB. • Storage-group name or ID from which to carve storage. • Volume name. • Volume disk image format (raw or qcow2).

	<ul style="list-style-type: none"> • Sector size of the volume (512 or 4096).
<code>--keep-volumes <i>volumes</i></code>	The idle volumes to be attached to the new 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.
<code>--storage-group <i>group</i></code>	The storage group from which to carve the VM volumes. If you do not specify this value, the storage group with the most free space is automatically selected. (If the storage group is configured with 4 kB sectors, ensure that the guest OS supports the 4 kB sector size.
<code>--vnc-keyboard-layout <i>keyboard</i></code>	The keyboard layout to assign to the VM. Valid options are English (<code>en-us</code>), German (<code>de</code>), or Japanese (<code>ja</code>). The default value is <code>en-us</code> .
<code>--no-auto-start</code>	If set, the VM is not started after the copy has finished.

Examples

Copy a VM named `vm:o2046` to a new vm called `new_vm_name` and retain all of the original settings.

```
$ avcli vm-copy --source-vm vm:o2046 --name new_vm_name
```

Copy a VM named `vm_source` to a new high availability VM named `vm_copy` with 2 CPUs and 1,024 MB of memory. Copy the boot volume `volume:o7652` to the storage group `storagegroup:o129` with the new name `vm_source_vol0_bootable_copy`, image type of `qcow2`, and sector size of 512 B. Copy the volume `volume:o7749` with the default values. Also, create a new volume named `vm_copy_add_new1` with size 20GB to `storagegroup:o1090` and set the image type to `qcow2` and sector size to 4 kB.

```
$ avcli vm-copy --source-vm vm_source --name vm_copy --cpu 2 -
-memory 1024 --availability ha --copy-volumes
volume:o7652,storagegroup:o129,vm_source_vol0_bootable_
copy,qcow2,512 volume:o7749 --add-volumes
20GB,storagegroup:o1090,vm_copy_add_new1,qcow2,4096
```

Copy a VM named `vm_source` to a new high availability VM named `new_vm_name` with 2 CPUs and 1,024 MB of memory. Copy the boot volume `boot_volume` to the Initial-Storage-Group with the new name `boot_volume_copy` and the image type of `raw`. Copy the volume `volume:o10158` with the default values. Create a new volume named `volume_new1` with size 20GB to `storagegroup:o71` and set the image type to `qcow2` and sector size to 4 kB. Attach two idle volumes, `volume_idle` and `volume:o19656`. In addition, configure network interfaces `network0` and `sharednetwork:o61`, set the default storage group to `storagegroup:o71`, set the keyboard layout as `de`, and prevent the VM from automatically starting upon creation.

```
$ avcli vm-copy --source-vm vm_source --name new_vm_name --cpu
2 --memory 1024 --availability ha --copy-volumes boot_
volume,Initial-Storage-Group,boot_volume_copy,raw
volume:o10158 --add-volumes 20GB,storagegroup:o71,volume_
new1,qcow2,4096 --keep-volumes volume_idle volume:o19656 --
interfaces network0 sharednetwork:o61 --storage-group
storagegroup:o71 --vnc-keyboard-layout de --no-auto-start
```

vm-create

Usage

```
avcli vm-create --name name --cpu number --memory memory --
cdrom cd-name | --kickstart template | --remote-file-path path
[--remote-type type] [--remote-username username] [--remote-
password password] --availability level [--interfaces
networks] [--disabled-interfaces networks] [--storage-group
group] [--vnc-keyboard-layout keyboard] --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> or <code>--remote-file-path</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> or <code>--remote-file-path</code> .
<code>--remote-file-path <i>path</i></code>	A remote ISO repository to use when booting the VM. You cannot specify this option with <code>--cdrom</code> or <code>--kickstart</code> .
<code>--remote-type <i>type</i></code>	The type of remote ISO repository specified in the <code>--</code>

	remote-file-path option. Valid options are samba or nfs.
--remote-username <i>user-name</i>	The user account to specify for access to the remote ISO repository specified in the --remote-file-path option. Required for samba repositories.
--remote-password <i>password</i>	The user password to specify for access to the remote ISO repository specified in the --remote-file-path option. Required for samba repositories.
--availability <i>level</i>	The availability level, high availability (ha) or fault tolerant (ft).
--interfaces <i>networks</i>	The list of networks to attach to the VM. Specify a network only once. The attached network must not be private.
--disabled-interfaces <i>networks</i>	A list of networks to attach to the VM, but which should not be enabled. Specify a network only once. The attached network must not be private.
--storage-group <i>group</i>	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.
--vnc-keyboard-layout <i>keyboard</i>	The keyboard layout to assign to the VM. Valid options are English (en-us), German (de), or Japanese (ja). The default value is en-us.
--volumes <i>volumes</i>	List of volumes to attach to this VM. A <i>volume</i> is made up of five components, separated by commas:

	<ul style="list-style-type: none"> • Size of the volume; required. By default, the volume size is specified in megabytes, but you can use standard qualifiers such as KB, MB, GB, and TB. • Storage-group name or ID from which to carve storage. • Volume name. • Volume disk image format (raw or qcow2). • Sector size of the volume (512 or 4096). <ul style="list-style-type: none"> ▪ The sector size is in bytes (B), 512 B by default. ▪ If the sector size of the storage group is 512 B, the sector size of the volume must be 512 B. ▪ If the sector size of the storage group is in 4096 B (4 kB), both 512 B and 4096 B are supported as the sector size of the volume. ▪ The boot volume must be presented with 512 B as the sector size.
<pre>--wait -w</pre>	<p>Wait for the command to complete.</p>

Examples

Create a HA VM named vm001 with a CPU, 512 MB of memory, a 1,024 MB volume, and that is attached to network0. Connect a remote ISO from an NFS share.

```
$ avcli vm-create --name vm001 --cpu 1 --memory 512 \
--remote-file-path 134.111.24.224:/developer/windows_7.iso
\
```

```
--remote-type nfs --availability ha --interfaces network0
--volumes 1024
```

Create a HA VM named `vm001` with a CPU, 512 MB of memory, a 1,024 MB volume, and that is attached to `network0`. Connect a remote ISO from a Samba share.

```
$ avcli vm-create --name vm001 --cpu 1 --memory 512 --
remote-file-path //134.111.31.228/Users/TEST/windows.iso -
-remote-type samba \
--remote-username TEST --remote-password abc123 --
availability ha --interfaces network0 --volumes 1024
```

Create a HA 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 \
--availability ha --interfaces network0 --volumes 1024
```

Create a FT 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 \
--availability ft --interfaces network0 --volumes 1024 \
--storage-group Pool-0001
```

Create a HA 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 \
--availability ha --interfaces network0 --volumes
1024,Pool-0001,vm001_vol0
```

Create a FT 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 \  
--availability ft --interfaces network0 network1 \  
--volumes 10GB,Pool-0001 50GB,Pool-0002
```

Create a HA VM based on a kickstart template.

```
$ avcli vm-create --name vm001 --cpu 1 --memory 512 --  
kickstart template:o81 \  
--availability ha --interfaces network0 --volumes 10GB
```

Create a HA VM with a CPU, 1024 MB of memory, a 20 GB qcow2-format bootable volume named `vm001_volu_boot`, a 1,024 MB data volume named `vm001_volu_data` with 4096 B sector size, and that is attached to `network0`.

```
$ avcli vm-create --name vm001 --cpu 1 --memory 1024 \  
--cdrom CentOS-6.7-x86_64-minimal.iso --availability ha --  
interfaces network0 \  
--volumes 20GB,Pool-0001,vm001_volu_boot,qcow2 1024,Pool-  
0002,\  
vm001_volu_data,qcow2,4096
```

vm-create-from-snapshot**Usage**

```
avcli vm-create --vm-snapshot-oid oid [--name name] [--cpu
number] [--memory memory] [--availability level] [--interfaces
networks] [--storage-group group] [--volumes volumes] [--
volume-prefix prefix] [--no-auto-start]
```

Description

The `vm-create-from-snapshot` command creates a new VM from a VM snapshot.

Options

<code>--vm-snapshot-oid <i>oid</i></code>	The vm snapshot OID from which to create the VM.
<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>--availability <i>level</i></code>	The availability level, high availability (ha) or fault tolerant (ft).
<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 from which to carve the VM volumes. If you do not specify this value, the storage group with the most free space is automatically selected. (If the storage group is configured with 4 kB sectors, ensure that the guest OS supports the 4 kB sector size.)
<code>--volumes <i>volumes</i></code>	Restrict the included volumes to those specified; oth-

	erwise, all volumes will be created. Specify volumes by configuration name or ID, boot volume first.
<code>--volume-prefix <i>prefix</i></code>	Adds the specified <i>prefix</i> to the beginning of the newly imported volume names to prevent conflicts with existing volumes on the system. For example, if a source volume is <code>ocean_boot</code> , and you specify <code>--volume-prefix new</code> , the imported volume would be <code>new-ocean_boot</code> .
<code>--no-auto-start</code>	If set, the VM is not started after creation has finished.

Examples

```
$ avcli vm-create-from-snapshot --vm-snapshot-oid
vmsnapshot:o41963 --name vm001
```

```
$ avcli vm-create-from-snapshot --vm-snapshot-oid
vmsnapshot:o41963 --name vm001 --availability ha --interfaces
network0 --volumes centos-boot centos-data --volume-prefix
minimal
```

```
$ avcli vm-create-from-snapshot --vm-snapshot-oid
vmsnapshot:o41963 --name vm001 --availability ha --interfaces
network0 net_143 --storage-group initial-group --volumes
centos-boot centos-data --volume-prefix minimal --no-auto-
start
```

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-export

Usage

```
avcli vm-export [--path pathname] [--volumes volumes] [--wait]
[--force] vm-name
```

Description

The `vm-export` command exports a VM in OVF/VHD format to the directory specified by *pathname*. The command first exports VHD files, followed by the OVF file. When the OVF file appears in *pathname*, the export is complete.



Note: Before you can start an export, you must mount a target Windows/CIFS or NFS share (from another system) in the everRun host operating system. For details, see ["Exporting a Virtual Machine from an everRun System"](#) on page 208.

Options

<code>--path <i>pathname</i></code>	A pathname relative to the export mount point to where the exported OVF is written.
<code>--volumes <i>name</i></code>	Restrict the exported volumes to those specified; otherwise, all volumes will be created. Specify volumes by configuration name or ID, boot volume first.
<code>--wait</code>	Wait for the export operation to complete. Specify this option to view the export progress.
<code>--force</code>	Force the VM to be exported, even if it is still running.
<i>vm-name</i>	Specify the name of the VM to be exported.

Examples

```
$ avcli vm-export --path exports/excalibur1 excalibur1
$ avcli vm-export --volumes volume:o1345 volume:o1389 --path
exports/excalibur1 excalibur1
```

vm-import**Usage**

```
avcli vm-import --archive filename.ova [--no-auto-start] [--cpu number] [--memory size] [--name vm-name] [--storage-groups groups] [--interfaces networks] [--remap-volumes] [--volumes volumes] [--volume-prefix prefix] [--data] [--force] [--silent] [--dry-run] [--throttle amount] [--use-https] [--protection-level level]
```

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>--remap-volumes</code>	First attempt to remap all volumes to the shared-mirrors as defined in the archive, after that the <code>--volumes</code> and <code>--storage-groups</code> rules are in effect.
<code>--volumes <i>volumes</i></code>	Import only these volumes. By default, all available volumes from the OVF are imported.
<code>--volume-prefix <i>prefix</i></code>	Adds the specified <i>prefix</i> to the beginning of the newly imported volume names to prevent conflicts with existing volumes on the system. For example, if a source volume is <code>ocean_boot</code> , and you specify <code>--volume-prefix new</code> , the imported volume would be <code>new-ocean_boot</code> .
<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 <i>amount</i></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.
<code>--protection-level <i>level</i></code>	The protection level to assign to the VM. Valid options are HA and FT (default).

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 --remap-volumes --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
$ avcli vm-import --name myVM --protection-level HA --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-network-disable**Usage**

```
avcli vm-network-disable --name name --nodes nodes --networks
networks
```

Description

The `vm-network-disable` command disables a VM's networks on selected nodes.

Options

<code>--name <i>name</i></code>	The name or ID of a VM.
<code>--nodes <i>nodes</i></code>	The name(s) or ID(s) of PM(s) on which to disable the networks.
<code>--networks <i>networks</i></code>	The names or IDs of networks to disable.

Examples

Disable net2 on node1 for a VM named MyVM.

```
$ avcli vm-network-disable --name MyVM --nodes node1 \
--networks net2
```

vm-network-enable

Usage

```
avcli vm-network-enable --name name --nodes nodes --networks  
networks
```

Description

The `vm-network-enable` command enables a VM's networks on selected nodes.

Options

<code>--name <i>name</i></code>	The name or ID of a VM.
<code>--nodes <i>nodes</i></code>	The name(s) or ID(s) of PM(s) on which to enable the networks.
<code>--networks <i>networks</i></code>	The names or IDs of networks to enable.

Examples

Enable net2 on node1 for a VM named MyVM.

```
$ avcli vm-network-enable --name MyVM --nodes node1 \  
--networks net2
```

vm-poweroff

Usage

```
avcli vm-poweroff [vm...] [--wait]
```

Description

The `vm-poweroff` command powers off the specified VMs.

Options

<i>vm</i>	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 [--new-name name] [--description "description"] [--cpu number] [--memory size] [--addVolumes volumes] [--deleteVolumes volumes] [--keepVolumes volumes] [--interfaces networks] [--disabled-interfaces networks] [--vnc-keyboard-layout keyboard] [--detach-boot-volume] [--attach-boot-volume name]
```

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>--new-name <i>name</i></code>	Specify a new name for the VM.
<code>--description "<i>description</i>"</code>	Specify a description for the VM.
<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>	The list of volumes to attach to this VM. A <i>volume</i> is made up of five components, separated by commas: <ul style="list-style-type: none"> • Size of the volume; required. By default, the volume size is specified in megabytes, but you can use standard qualifiers such as KB, MB, GB, and TB.

	<ul style="list-style-type: none"> • Storage-group name or ID from which to carve storage. • Volume name. • Volume disk image format (raw or qcow2). • Sector size of the volume (512 or 4096). <ul style="list-style-type: none"> ▪ The sector size is in bytes (B), 512 B by default. ▪ If the sector size of the storage group is 512 B, the sector size of the volume must be 512 B. ▪ If the sector size of the storage group is in 4096 B (4 kB), both 512 B and 4096 B are supported as the sector size of the volume. ▪ The boot volume must be presented with 512 B as the sector size.
<p><code>--deleteVolumes <i>volumes</i></code></p>	<p>The list of volume(s) that are currently attached to the specified VM to be deleted. Specify a volume either by name or ID.</p>
<p><code>--keepVolumes <i>volumes</i></code></p>	<p>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.</p>
<p><code>--interfaces <i>networks</i></code></p>	<p>The list of networks to attach to the VM. Specify a network only once. The attached network must not be private.</p>
<p><code>--disabled-interfaces <i>networks</i></code></p>	<p>A list of networks to attach to the VM, but which should not be enabled. Specify a network only once. The attached network must not be private.</p>

<code>--vnc-keyboard-layout</code> <i>keyboard</i>	The keyboard layout to assign to the VM. Valid options are English (en-us), German (de), or Japanese (ja). The default value is en-us.
<code>--detach-boot-volume</code>	Detaches the VM's boot volume.
<code>--attach-boot-volume</code> <i>name</i>	Specify the name of a new boot volume for this VM. If the VM already has an attached boot volume, you must also specify <code>--detach-boot-volume</code> ; otherwise, the command fails.

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
```

Reprovision VM `vm1` with the same parameters from the previous example. Also rename the VM `vm2` and add a description.

```
$ avcli vm-reprovision --cpu 3 --memory 3359 --addVolumes  
2500,storagegroup:o54,data-1-7,qcow2 --deleteVolumes  
volume:o1043 --keepVolumes volume:o1 volume:o2 volume:o4 -  
-interfaces sharednetwork:o129 sharednetwork:o130 --name  
vm1 --new-name vm2 --description "This is the vm  
description"
```

Reprovision VM `vm001` with two CPUs, 2048 MB of memory, one new data volume `vm001_data1` in `qcow2` format with 4 kB sector size, and keep volume `o7517`.

```
$ avcli vm-reprovision --cpu 2 --memory 2048 --addVolumes  
20GB,storagegroup:o1090,vm001_data1,qcow2,4096 --  
keepVolumes volume:o7517 --name vm001
```

Swap two VMs' boot disks.

Detach boot volume:

```
$ avcli vm-reprovision --detach-boot-volume --name p56xen4
```

Switch boot volume:

```
$ avcli vm-reprovision --detach-boot-volume --attach-boot-  
volume boot-p56xen4 --name p56xen8
```

Attach a detached boot volume to a different VM:

```
$ avcli vm-reprovision --attach-boot-volume boot-p56xen8 -  
-name p56xen4
```

vm-restore

Usage

```
avcli vm-restore --archive filename.ova [--no-auto-start][--cpu number][--memory size][--name vm-name][--storage-groups groups][--interfaces networks][--volume-prefix prefix][--data] [--silent][--dry-run] [--throttle][--use-https]
```

Description

The `vm-restore` command restores a VM from an OVA or OVF file.

Options

<code>--archive filename.ova</code>	The OVA or OVF file archive to restore.
<code>--no-auto-start</code>	Do not start the VM after the restore has finished.
<code>--cpu number</code>	The number of CPUs to assign to the VM. This defaults to the value in the archive.
<code>--memory size</code>	The size of memory, in megabytes, to assign to the VM. This defaults to the value in the archive.
<code>--name vm-name</code>	The name to assign to the VM. This defaults to the value in the archive.
<code>--storage-groups groups</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 networks</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>--volume-prefix prefix</code>	Adds the specified <i>prefix</i> to the beginning of the

	<p>newly imported volume names to prevent conflicts with existing volumes on the system. For example, if a source volume is <code>ocean_boot</code>, and you specify <code>--volume-prefix new</code>, the imported volume would be <code>new-ocean_boot</code>.</p>
<code>--data</code>	Restore data only for the specified volumes.
<code>--silent</code>	Suppress output.
<code>--dry-run</code>	Show the interface to the shared network and volume-to-storage-group assignments without actually restoring a VM.
<code>--throttle</code>	<p>Slow down the operation. Valid values are:</p> <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-restore --archive vm1.ova
```

```
$ avcli vm-restore --archive vm1/vm1.ovf
```

```
$ avcli vm-restore --name myVM --throttle low --archive
vm1.ovf
```

```
$ avcli vm-restore --cpu 2 --memory 1024 --archive vm1.ovf
```

```
$ avcli vm-restore --interfaces network0 network1 --archive  
vm1.ovf
```

```
$ avcli vm-restore --storage-groups sm-0000 sm-0001 --archive  
vm1.ovf
```

```
$ avcli vm-restore --data vol1 vol3 --archive vm1.ovf
```

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-snapshot-create

Usage

```
avcli vm-snapshot-create [--volumes | --no-data] [--description] [--desire] [--require] vm-name
```

Description

The `vm-snapshot-create` command creates a VM snapshot.

Two snapshot consistency levels are supported:

- *Crash consistency*: The restored data is in the same state that it would have been if the system had crashed at the exact moment that the snapshot was taken. A crash-consistent snapshot does not capture the contents of memory or any pending I/O operations.
- *Application consistency*: Before the snapshot is taken, cooperating applications are briefly frozen so that transactions are complete, buffers flushed, files closed, and so on. This ensures that cooperating applications start from a consistent state. This is the highest level of consistency.

Options

<code>--volumes --no-data</code>	The names of volumes to include in the snapshot. By default, all volumes are included in the snapshot unless you specify <code>--volumes</code> with individual volume names or you specify <code>--no-data</code> . If you specify <code>--no-data</code> , no volumes are included in the snapshot. These arguments are mutually exclusive.
<code>--description</code>	The user-specified description for this snapshot.
<code>--desire</code>	The highest consistency level to attempt in order to declare the snapshot a success. If this attempt fails, attempts are made at successively lower levels (but no lower than the level specified by <code>--require</code>). Values are <code>crash</code> and <code>application</code> (the default value).
<code>--require</code>	The minimum consistency level needed to declare the snapshot a

	success. Values are <code>crash</code> and <code>application</code> (the default value).
<i>vm-name</i>	The VM's ID.

Examples

```
$ avcli vm-snapshot-create --volumes volume:o100  
volume:o101 vm1
```

vm-snapshot-create-disable

Usage

```
avcli vm-snapshot-create-disable
```

Description

The `vm-snapshot-create-disable` command disables the system's ability to create snapshots. By default, the system's ability to create snapshots is enabled. Only users with the role **Administrator** (the group `admin`) can execute this command. Users with the role **Platform Admin** (the group `platform_admin`) or **Read Only** (the group `read-only`) cannot execute this command.

Examples

```
$ avcli -H localhost -u admin -p password vm-snapshot-create-  
disable
```

vm-snapshot-create-enable

Usage

```
avcli vm-snapshot-create-enable
```

Description

The `vm-snapshot-create-enable` command enables the system's ability to create snapshots. By default, the system's ability to create snapshots is enabled. Only users with the role **Administrator** (the group `admin`) can execute this command. Users with the role **Platform Admin** (the group `platform_admin`) or **Read Only** (the group `read-only`) cannot execute this command.

Examples

```
$ avcli -H localhost -u admin -p password vm-snapshot-create-  
enable
```

vm-snapshot-delete**Usage**

```
avcli vm-snapshot-delete snapshot...
```

Description

The `vm-snapshot-delete` command deletes the specified snapshots.

Options

<i>snapshot</i>	One or more snapshots of the VM. Specify a snapshot by ID.
-----------------	--

Examples

```
$ avcli vm-snapshot-delete vmsnapshot:o100 vmsnapshot:o101
```

vm-snapshot-export

Usage

```
avcli vm-snapshot-export [--wait][--silent][--volumes volumes]
--path pathname snapshot
```

Description

The `vm-snapshot-export` command exports a VM in OVF/VHD format to the directory specified by *pathname*. The command first exports VHD files, followed by the OVF file. When the OVF file appears in *pathname*, the export is complete.



Note: Before you can start an export, you must mount a target Windows/CIFS or NFS share (from another system) in the everRun host operating system. For details, see ["Exporting a Snapshot" on page 256](#).

Options

<code>--wait</code>	Wait for the export operation to complete. Specify this option to view the export progress.
<code>--silent</code>	Suppress progress output.
<code>--volumes <i>volumes</i></code>	Restrict the exported volumes to those specified; otherwise, all volumes will be created. Specify volumes by configuration name or ID, boot volume first.
<code>--path <i>pathname</i></code>	A pathname relative to the export mount point to where the exported OVF is written.
<i>snapshot</i>	The name of the snapshot to export.

Examples

Export a snapshot with all captured volumes:

```
$ avcli vm-snapshot-export --path exports/ex1 ex1
```

Export a snapshot with just one captured volume:

```
$ avcli vm-snapshot-export --volumes boot-ex1 --path  
exports/ex1 ex1
```

vm-snapshot-info

Usage

```
avcli vm-snapshot-info [snapshot...]
```

Description

The `vm-snapshot-info` command displays information about all snapshots, or optionally, only about those specified.

Options

<i>snapshot</i>	One or more snapshots of the VM. Specify a snapshot either by name or ID.
-----------------	---

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
```

vm-volume-disable

Usage

```
avcli vm-volume-disable --name name --nodes nodes --volumes  
volumes
```

Description

The `vm-volume-disable` command disables a VM's volumes on selected nodes.

Options

<code>--name <i>name</i></code>	The name or ID of a VM.
<code>--nodes <i>nodes</i></code>	The name(s) or ID(s) of PM(s) on which to disable the volumes.
<code>--volumes <i>volumes</i></code>	The names or IDs of volumes to disable.

Examples

Disable `abba1-data` and `volume:o2249` on `node1` for a VM named `MyVM`.

```
$ avcli vm-volume-disable --name MyVM --nodes node1 \  
--volumes abba1-data volume:o2249
```

vm-volume-enable**Usage**

```
avcli vm-volume-enable --name name --nodes nodes --volumes
volumes
```

Description

The `vm-volume-enable` command enables a VM's volumes on selected nodes.

Options

<code>--name <i>name</i></code>	The name or ID of a VM.
<code>--nodes <i>nodes</i></code>	The name(s) or ID(s) of PM(s) on which to enable the volumes.
<code>--volumes <i>volumes</i></code>	The names or IDs of volumes to enable.

Examples

Enable volume:o2249 on node1 for a VM named MyVM.

```
$ avcli vm-volume-enable --name MyVM --nodes node1 \
--volumes volume:o2249
```

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.
---------------	--

volume-resize

Usage

```
avcli volume-resize --new-size size volume
```

Description

The `volume-resize` command resizes a volume. The image container (also known as a *volume container*) must be large enough to allow this. You must stop the VM before specifying this command.

Options

<code>--new-size <i>size</i></code>	The new volume size. By default, <i>size</i> is in megabytes, but you can specify standard qualifiers (for example, KB, K, MB, M, GB, or G).
<code><i>volume</i></code>	The volume to be resized.

Examples

```
# avcli volume-resize --new-size 79G boot-airplane1
```


12

Chapter 12: System Reference Information

See the following topics for reference information

- ["Compatible Guest Operating Systems" on page 474](#)
- ["Physical Machine System Requirements" on page 476](#)
- ["Important Physical Machine and Virtual Machine Considerations" on page 478](#)
- ["Accessing Knowledge Base Articles" on page 481](#)

Compatible Guest Operating Systems

The following are compatible as guest operating systems for virtual machines (VMs) running on everRun systems.

Operating System	Version
CentOS 7	CentOS 7.0, 7.1, 7.2, 7.3 (all 64-bit)
CentOS 6	CentOS 6.4, 6.5, 6.6, 6.7, 6.8, 6.9 (all 64-bit)
Microsoft Windows Server 2016 (Essentials, Standard, Data-center)	64-bit
Microsoft Windows Server 2012 (Foundation, Essentials, Stand-	64-bit, 64-bit R2

Operating System	Version
ard, Datacenter)	
Microsoft Windows Small Business Server 2011 (Standard, Essential, Premium Add-On)	64-bit
Microsoft Windows Server 2008 (Web, Small Business, Standard, Enterprise, Datacenter)	32-bit, 64-bit R2 only
Microsoft Windows Server 2003 (Enterprise)	32-bit R2 SP2 ¹
Microsoft Windows 10 Desktop	64-bit
Microsoft Windows 8.1 Desktop Enterprise)	64-bit
Microsoft Windows 8 Desktop (Enterprise)	64-bit
Microsoft Windows 7 Desktop	32-bit, 64-bit
Oracle Linux 7.3 Red Hat Compatible (RHC) Kernel	x86-64
Oracle Linux 6.8 RHC Kernel	x86-64
Red Hat Enterprise Linux 7 (Workstation, Server)	Red Hat 7.0, 7.1, 7.2, 7.3 (all 64-bit)
Red Hat Enterprise Linux 6 (Workstation, Server)	Red Hat 6.4, 6.5, 6.6, 6.7, 6.8, 6.9 (all 64-bit)
SUSE Linux Enterprise Server	SLES 11 SP3 64-bit SLES 11 SP4 64-bit SLES 12 SP2 64-bit
Ubuntu	12.04, 13.10, 14.04, 16.04 (all 64-bit)

¹For specific installation and migration procedures, see ["Creating a New Windows Server 2003 Virtual Machine" on page 147](#) and ["Migrating a Windows Server 2003 VM to an everRun 7.x System" on page 179](#).

Physical Machine System Requirements

The following table presents the minimum and maximum capacities of the listed devices for physical machines running on everRun systems.

Physical Device	Minimum	Maximum Tested	Architected	Notes
Processors/CPUs:				
Intel® Xeon® Gold 6XXX, 5XXX				
Intel Xeon Silver 4XXX				
Intel Xeon Bronze 3XXX				
Intel Xeon E5- 2XXX v4				
Intel Xeon E5- 2XXX v3	1	2	No Practical Limit	Maximum of 2 phys- ical CPUs supported in one physical machine (PM) of an everRun system.
Intel Xeon E5- 2XXX v2				
Intel Xeon E5- 2XXX				
Intel Xeon E5- 1XXX v4				
Intel Xeon E5- 1XXX v3				
Intel Xeon E5- 1XXX v2				

Physical Device	Minimum	Maximum Tested	Architected	Notes
Intel Xeon E5-1XXX				
Intel Xeon E3-1XXX v6				
Intel Xeon E3-1XXX v5				
Intel Xeon E3-1XXX v4				
Intel Xeon E3-1XXX v3				
Intel Xeon E3-1XXX v2				
Intel Xeon E3-1XXX				
Number CPU Sockets per Physical Machine	1	2	No Practical Limit	Maximum of 2 physical sockets supported in one PM of an everRun system.
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	

Physical Device	Minimum	Maximum Tested	Architected	Notes
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 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 476](#)
- ["Virtual Machine Recommendations and Limits" on page 478](#)
- ["Combined Virtual Machine Maximums" on page 480](#)
- ["Important Considerations" on page 480](#)

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 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 system resources to the maximum.

The following table lists everRun system VM limits.

Item	Limits
Maximum VCPUs per FT VM	20
Maximum VCPUs per HA VM	20
Maximum Memory per FT VM	256 GB
Maximum Memory per HA VM	256 GB
Maximum Availability Links per VM	2
Maximum Virtual Networks per VM	20
Maximum Storage Volumes per VM	12
Guest Volume Size	Tested up to 2.2 TB. Maximum supported size is 16 TB.
Max. Snapshots per VM	16 (72 total per system)

Combined Virtual Machine Maximums

The following table presents the combined maximums of virtual machines (VMs) and virtual NICs that can run on everRun systems.

Virtual Device	Maximum
Total FT VMs	8
Total VMs (FT and HA combined)	28
Total Virtual Network Interface cards (NICs)	20

Important Considerations

Note the following important considerations.

Feature	Comment
everRunSystem Disk	<p>Recommended minimum configuration for physical machines:</p> <ul style="list-style-type: none"> • One logical volume protected by RAID1, RAID 5, RAID 6, or RAID 10 or • 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	<p>USB CD/DVD drives are supported on all platforms for everRun installation.</p>
Direct-Attach Tape Drives	<p>Access to direct-attach tape drives by the guests is not supported. Stratus recommends using network-attach tape drives.</p>
Console Connectivity	<p>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.</p>
SSD Support	<p>everRun supports solid state drives according to the storage controller vendor specifications.</p>
System Management	<p>everRun system management does not work at run-level 5.</p>

Accessing Knowledge Base Articles

The **Stratus Customer Service Portal** provides a searchable **Knowledge Base** with technical articles about all Stratus products, including everRun. In some cases, the online Help directly references these Knowledge Base articles (for example, KB-*nnnn*). You can access the Customer Service Portal and

Knowledge Base by using your existing service portal credentials, or by creating a new user account, as follows.

To access the Knowledge Base

1. Log on to the **Stratus Customer Service Portal** at <https://support.stratus.com>.

If needed, create a new account as follows:

- a. Click **Register Account**.
- b. Enter your company email address and contact info, and click **Register**.

Your company email address must include a domain name (for example, stratus.com) for a company that is a registered customer of Stratus.

- c. Click the link in the email that you receive from Stratus.
- d. Enter a new password and finish configuring your account.

If you need assistance creating an account, contact your authorized Stratus service representative.

2. In the service portal, click **Knowledge Base** in the left pane.
3. In the **Keyword Search** box, enter keywords associated with the information you need, and then click **Search**.

To search for an article by its KB-*nnnn* number, click **Advanced Search**. Next to **Search by ID**, type the article ID number (*nnnn*) and click **Display**.

Related Topics

["Supporting Documents" on page 287](#)

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 system to use SNMP, see ["Configuring SNMP Settings" on page 89](#).

For information on using the `snmptable` command to obtain information about the system, specifically information about alerts, audit logs, nodes, VMs, and volumes, see ["Obtaining System Information with snmptable" on page 484](#).

You can download a copy of the MIB file from the **Drivers and Tools** section of the **Downloads** page at <http://www.stratus.com/go/support/everrun>.

Obtaining System Information with snmptable

You can issue the `snmptable` command to obtain information about the system, specifically information about alerts, audit logs, nodes, VMs, and volumes.

To display alert information

To display information about alerts, issue the following command:

```
snmptable -v2c -m+/usr/smd/STRATUS-EVERRUN-MIB.txt -c  
public localhost everRunAlertTable
```

The command output displays the following:

Field	Description												
everRunAlertIndex	The alert number.												
everRunAlertSeverity	<p>The alert severity (see everRunAlertSeverityNum for numerical value).</p> <p>Values are:</p> <table border="0"> <tr> <td>clear</td> <td></td> </tr> <tr> <td>informational</td> <td></td> </tr> <tr> <td>minor</td> <td></td> </tr> <tr> <td>major</td> <td></td> </tr> <tr> <td>serious</td> <td></td> </tr> <tr> <td>critical</td> <td></td> </tr> </table>	clear		informational		minor		major		serious		critical	
clear													
informational													
minor													
major													
serious													
critical													
everRunAlertType	<p>The type of alert. Examples are:</p> <ul style="list-style-type: none"> • node_singleSystemDisk • Node Maintenance • The Unit is not well balanced 												
everRunAlertSource	<p>The source of the alert. Examples are:</p> <ul style="list-style-type: none"> • node0 or node1 • everRun system network name • network host name 												
everRunAlertDateTime	<p>The date and time of the alert, in the format <i>yyyy-mm-dd hh:mm:ss</i>, where <i>yyyy</i> is year, <i>mm</i> is month, <i>dd</i> is date, <i>hh</i> is hour, <i>mm</i> is minute, and <i>ss</i> is second (for example, 2017-11-03 23:49:45).</p>												
everRunAlertCallHomeS	If true, Call Home was sent; if false, it was not sent												

Field	Description												
ent													
everRunAlertEAlertSent	If true, e-Alert was sent; if false, it was not sent												
everRunAlertSNMPTrapSent	If true, SNMP trap was sent; if false, it was not sent												
everRunAlertInformation	<p>Information about the alert. Examples are:</p> <ul style="list-style-type: none"> • Node node1 is in maintenance • node0 has a single system disk: Policy assumes this disk is redundant - if not, please add another internal disk • BUSINESS network net_728 is reporting a degraded link condition • The unit is not well load balanced 												
everRunAlertSNMPTrapOID	SNMP trap object identifier (OID) (for example, COMPANY-MIB::nodeSingleSystemDisk)												
everRunAlertSeverityNumber	<p>everRunAlertSeverity number. Values are:</p> <table border="0"> <tr> <td style="padding-right: 20px;">0</td> <td>Clear</td> <td></td> </tr> <tr> <td>1</td> <td>Informational</td> <td></td> </tr> <tr> <td>2</td> <td>Minor</td> <td></td> </tr> <tr> <td>3</td> <td>Major</td> <td></td> </tr> </table>	0	Clear		1	Informational		2	Minor		3	Major	
0	Clear												
1	Informational												
2	Minor												
3	Major												

Field	Description
	4 Serious 
	5 Critical 

To display audit log information

To display information about audit logs, issue the following command:

```
snmptable -v2c -m+/usr/smd/STRATUS-EVERRUN-MIB.txt -c
public localhost everRunAuditTable
```

The command output displays the following:

Field	Description
everRunAuditIndex	An incrementing number (1, 2, etc.) to indicate the audit log whose information is displayed..
everRunAuditDateTime	The date and time that the audit was generated, in the format <i>yyyy-mm-dd hh:mm:ss</i> , where <i>yyyy</i> is year, <i>mm</i> is month, <i>dd</i> is date, <i>hh</i> is hour, <i>mm</i> is minute, and <i>ss</i> is second (for example, 2017-11-03 23:49:45).
everRunAuditUsername	The name of the user that caused the audit to be generated (for example, <i>audit</i> or <i>admin</i>).
everRunAuditOriginatingHost	The IP address of the host that originated the audit.
everRunAuditAction	A description of the action being audited. Examples are: <ul style="list-style-type: none"> • "Login user \"audit" • "Start virtual machine \"manager1" • "Remove all cleared alert"

To display node information

To display node information, issue the following command:

```
snmptable -v2c -m+/usr/smd/STRATUS-EVERRUN-MIB.txt -c
public localhost everRunNodeTable
```

The command output displays the following:

Field	Description
everRunNodeIndex	A number (typically 1 or 2) to indicate the node whose information is displayed.
everRunNodeId	The host ID of the node (for example, host : o34).
everRunNodeDisplayName	The node name, node0 or node1.
everRunNodeIsPrimary	If true, the node is primary. If false, the node is secondary.
everRunNodeStateNum	<p>Node state is:</p> <ul style="list-style-type: none"> 0 Normal (✓) 1 Warning (⚠) 2 Busy (🔄) 3 Broken (✖) 4 Maintenance (🛠)
everRunNodeActivityNum	<p>Node activity is:</p> <ul style="list-style-type: none"> 0 Imaging 1 Booting 2 Running 3 Stopping 4 Rebooting

Field	Description
	5 Powered off
	6 Failed
	7 Firmware updating
	8 Lost
	9 Exiled
	10 Unreachable
	11 Proto (initializing)
	12 Evacuating

To display VM information

To display VM information, issue the following command:

```
snmptable -v2c -m+/usr/smd/STRATUS-EVERRUN-MIB.txt -c
public localhost everRunVMTable
```

The command output displays the following:

Field	Description
everRunVMIndex	An incrementing number (1, 2, etc.) to indicate the VM whose information is displayed.
everRunVMId	The VM ID (for example, vm:01467).
everRunVMDisplayName	The VM name (for example, MyVM).
everRunVMRunningNode	The node on which the VM is running, node0 or node1.
everRunVMAvailability	The VM availability, HA (High Availability) or FT (Fault Tolerant).

Field	Description
everRunVMStateNum	VM state is: 0 Normal (✔) 1 Warning (⚠) 2 Busy or synchronizing (🔄) 3 Broken or blacklisted (✖)
everRunVMActivityNum	VM activity is: 0 Installing 1 Booting 2 Running 3 Moving 4 Stopping 5 Stopped 6 Exporting 7 Taking snapshot 8 Paused 9 Loading 10 Crashing 11 Crashed 12 Dumping 13 Waiting

To display volume information

To display volume information, issue the following command:

```
snmptable -v2c -m+/usr/smd/STRATUS-EVERRUN-MIB.txt -c
public localhost everRunVolumeTable
```

The command output displays the following:

Field	Description
everRunVolumeIndex	An incrementing number (1, 2, etc.) to indicate the volume whose information is displayed.
everRunVolumeId	The volume ID (for example, volume : o588).
everRunVolumeDisplayN ame	The volume name (for example, root).
everRunVolumeSyncPerc entage	The percentage of the volume that is synchronized.
everRunVolumeStorageG roup	The storage group that the volume is part of.
everRunVolumeUsedBy	The name of the VM or host that is using the volume (for example, MyVM); none indicates that the volume is not being used.
everRunVolumeStateNum	<p>Volume state is:</p> <p>0 Normal (✓)</p> <p>1 Warning (⚠)</p> <p>2 Busy or synchronizing (🔄)</p> <p>3 Broken (✖)</p>