

everRun User's Guide



Notice

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

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

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

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

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 change, 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 Ireland, Ltd. The Stratus Technologies logo, the Stratus 24 x 7 logo, and Automated Uptime are trademarks of Stratus Technologies Ireland, 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, vSphere, and ESXi are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions.

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

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

Mozilla and Firefox are registered trademarks of the Mozilla Foundation.

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

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 7.5.1.0 Publication Date: Thursday, April 29, 2021 Stratus Technologies, Inc. 5 Mill and Main Place, Suite 500 Maynard, Massachusetts 01754-2660 © 2020 Stratus Technologies Ireland, Ltd. All rights reserved.

Table of Contents

Part 1: everRun User's Guide	
Chapter 1: Introduction to everRun	Systems 1
everRun System Overview	
everRun System Description	
Physical Machines and Virtual Ma	chines 2
Administrative Operations	
Alerts	
Remote Support	
Lights-Out Management	
Third-party Management Tools	
Modes of Operation	
High Availability Operation	
Fault Tolerant Operation	
SplitSite Configurations	
SplitSite and Quorum Service.	
Quorum Servers	
everRun Storage Architecture	
Logical Disks and Physical Disks	
Storage Groups	
Sizing Volume Containers	
Network Architecture	
A-Link and Private Networks	
Business and Management Netw	orks13
Network Segmentation Fault Dete	ection and Remediation14
System Usage Restrictions	
QEMU	
Accessing the Host Operating Sy	stem15
Chapter 2: Getting Started	
Planning	
System Requirements Overview	
System Hardware	
Supported Servers	
RAM	

Disk Space	
Network	18
IP Addresses	19
Ports	
Storage Requirements	
Memory Requirements	21
General Network Requirements and Configurations	
Requirements	22
Recommended Configurations	
Business and Management Network Requirements	24
A-Link and Private Network Requirements	
everRun Availability Console Requirements	26
Compatible Internet Browsers	27
Power Requirements and Considerations	27
Software Installation	27
Site and System Preparation	27
Connecting Power	
UPS (Optional)	28
Obtaining everRun Software	
Obtaining the ISO Image	
Final Step	
Creating Bootable USB Media	
Configuring Settings in the Firmware Setup Utility	
Required Settings	34
Recommended Settings	
Installing everRun Software	
Connecting Ethernet Cables	
Installation Options	
Installing Software on the First PM	
Mapping Your Keyboard	
To configure your keyboard layout during installation:	45
To configure your keyboard layout after installation:	
Recording the Management IP Address	45
Installing Software on the Second PM	46
Post-Installation Tasks	48

Obtaining System IP Information	49
Logging On to the everRun Availability Console for the First Time	49
Connecting Additional Networks	51
Chapter 3: Using the everRun Availability Console	53
The everRun Availability Console	54
Logging On to the everRun Availability Console	55
The Dashboard Page	56
Resolving Outstanding Alerts on the Dashboard	56
The System Page	
Rebooting the System	58
Shutting Down the System	59
The Preferences Page	60
Specifying Owner Information	63
Managing the Product License	63
Configuring IP Settings	67
Configuring Quorum Servers	69
Configuring Date and Time	71
Configuring System Resources	72
Configuring Active Directory	73
Configuring the Migration Policy	74
Configuring Secure Connections	75
Configuring the Host Inactivity Logout	79
Disabling and Enabling Snapshots	80
Configuring VM Devices	
Managing IPtables	81
Configuring e-Alerts	
Configuring SNMP Settings	
Configuring Remote Support Settings	
Configuring Internet Proxy Settings	90
The Alerts History Page	91
The Audit Logs Page	92
The Support Logs Page	92
Creating a Diagnostic File	
Uploading a Diagnostic File to Customer Support	94
Deleting a Diagnostic File	95

The Physical Machines Page	95
Physical Machine Actions	96
Physical Machine States and Activities	97
The Virtual Machines Page	
Virtual Machine Actions	100
Virtual Machine States and Activities	103
The Snapshots Page	
The Volumes Page	105
The Storage Groups Page	106
The Networks Page	107
Fixing a Network Connection	108
Setting the MTU	109
The Virtual CDs Page	
The Upgrade Kits Page	110
Creating a USB Medium with System Software	110
The Users & Groups Page	111
Managing Local User Accounts	112
Managing Domain User Accounts	113
Chapter 4: Upgrading everRun Software	115
Upgrading everRun Software Using an Upgrade Kit	115
Upgrading everRun From Release 7.3.4.0 with a DVD	119
Chapter 5: Managing Logical Disks	121
Logical Disk Management	121
Responding to a Failed Logical Disk	
Activating a New Logical Disk	124
Creating a New Storage Group	
Deleting a Storage Group	125
Assigning a Logical Disk to a Storage Group	
Chapter 6: Managing Physical Machines	129
Maintenance Mode	129
Rebooting a Physical Machine	131
Shutting Down a Physical Machine	
Load Balancing	133
Modes of Operation	
Troubleshooting Physical Machines	

Recovering a Failed Physical Machine	134
Chapter 7: Managing Virtual Machines	141
Planning Virtual Machine Resources	142
Planning Virtual Machine vCPUs	143
Planning Virtual Machine Memory	144
Planning Virtual Machine Storage	145
Planning Virtual Machine Networks	
Creating and Migrating Virtual Machines	148
Creating a New Virtual Machine	
Creating a New Windows Server 2003 Virtual Machine	
Copying a Virtual Machine	
Migrating a Physical Machine or Virtual Machine to a System	157
Migrating From Avance or everRun MX Systems	
Planning to Migrate from an everRun MX System	
Platform Requirements	171
Planned Outage	171
Guest Operating System Support	
Network Preparation	171
Storage Considerations	
Quorum Support	172
Installing everRun Software	173
Migrating Virtual Machines	173
Converting an everRun MX System to an everRun 7.x System	173
Planning to Migrate from an Avance Unit	
Platform Requirements	
Planned Outage	
Guest Operating System Support	
Network Preparation	
Storage Considerations	
Installing everRun Software	182
Migrating Virtual Machines	
Converting an Avance Unit to an everRun 7.x System	
Migrating a Windows Server 2003 VM to an everRun 7.x System	
Importing an OVF File from an everRun MX System	191
Importing an OVF File from an Avance System	199

Importing an OVF File
Replacing/Restoring a Virtual Machine from an OVF File
Exporting a Virtual Machine
Mounting a USB Device or Network-mounted Folder on the everRun System
Managing Windows Drive Labels
Configuring Windows-based Virtual Machines
Updating the VirtIO Drivers (Windows-based VMs)
Creating and Initializing a Disk (Windows-based VMs)230
Installing Applications (Windows-based VMs)231
Installing the QEMU Guest Agent for Application-Consistent Snapshots (Windows-based VMs)231
Configuring Linux-based Virtual Machines
Creating and Initializing a Disk (Linux-based VMs)
Installing Applications (Linux-based VMs)
Installing the QEMU Guest Agent for Application-Consistent Snapshots (Linux-based VMs)235
Managing the Operation of a Virtual Machine
Starting a Virtual Machine
Shutting Down a Virtual Machine
Powering Off a Virtual Machine
Opening a Virtual Machine Console Session
Renaming a Virtual Machine
Removing a Virtual Machine
Managing Virtual Machine Resources
Reprovisioning Virtual Machine Resources
Creating a Volume in a Virtual Machine
Attaching a Volume to a Virtual Machine
Detaching a Volume from a Virtual Machine
Removing a Volume from a Virtual Machine
Renaming a Volume on the everRun System
Expanding a Volume Container on the everRun System
Expanding a Volume on the everRun System
Recovering Virtual Machine Resources
Enabling and Disabling VM Components
Managing Virtual CDs
Creating a Virtual CD
Inserting a Virtual CD

Ejecting a Virtual CD	
Booting from a Virtual CD	
Renaming a Virtual CD	
Removing a Virtual CD	
Managing Snapshots	
Creating a Snapshot	
Creating a Virtual Machine from a Snapshot	
Exporting a Snapshot	270
Removing a Snapshot	
Advanced Topics (Virtual Machines)	275
Assigning a Specific MAC Address to a Virtual Machine	
Selecting a Preferred PM for a Virtual Machine	277
Forcing a VM to Boot	
Changing the Protection Level for a Virtual Machine (HA or FT)	
Configuring the Boot Sequence for Virtual Machines	
Resetting MTBF for a Failed Virtual Machine	
Locating a Dump File in a Virtual Machine	
Attaching a USB Device to a Virtual Machine	
Chapter 8: Maintaining Physical Machines	
Physical Machine Hardware Maintenance Restrictions	
Adding or Replacing Hot-Swappable Components	
Adding or Replacing Components That Are Not Hot-Swappable	
Adding a New NIC	
Replacing Physical Machines, Motherboards, NICs, or RAID Controllers	
Upgrading Both Physical Machines In a Running System	
Part 2: Supporting Documents	
Chapter 9: everRun Release 7.5.1.0 Release Notes	
Accessing Stratus Knowledge Base Articles	
Bug Fixes	
Fixed in everRun Release 7.5.1.0	
New Features and Enhancements	
New in everRun Release 7.5.1.0	306
New in everRun Release 7.5.0.5	
Important Considerations	
Upgrading from everRun Release 7.4.x.x	

Upgrading from everRun Release 7.3.4.x or Earlier	308
Re-activating the everRun Product License After Replacing a PM	310
Other Important Considerations	
Known Issues	310
P2V/V2V Migration May Fail when Migrating a Windows-based VM	310
VM Network Remains Offline After Using P2V/V2V to Migrate from a VMware Source	311
VM Network Remains Offline After Importing a Windows-based VM	311
Ensure that Both Nodes Are Running Before Removing a VM	311
Invalid VM Names Prevent Some VM Operations from Executing	311
Guest OS May Fail to Display a VCD When Inserted	311
Some Browsers Unable to Connect a VNC When Using https	312
Cannot Mount USB Media with exFAT File System	312
Recovering a Missing or Failed Logical Disk with the Repair Button May Be Slow	312
Removing Snapshots Temporarily Prevents Some VM Operations	313
Creation of Snapshots Results in Volume Conversion from RAW to QCOW3 Format	313
Moving an everRun System to a Different Subnet	314
Enabling Log Files for the Windows QEMU Guest Agent May Cause VM Timeouts when T	ak-
ing Snapshots	314
Reboot Required when Changing Node IP Address or Netmask Network Settings	314
Changing the MTU of Business Networks Causes VM Migration and Possibly Node Failov	er314
Documentation Updates	314
Updates in everRun Release 7.5.0.5 Documentation	314
Getting Help	315
Chapter 10: everRun Command Line Interface Reference	316
AVCLI Command Overview	316
Prerequisites	317
Installing the Linux Client	317
Installing the Windows Client	318
Using AVCLI	319
Executing a Command	319
Using AVCLI Help	320
Listing All Commands	320
Displaying Help for a Specific Command	321
AVCLI Error Status	321
XML Encapsulated Errors	321

Error Checking	322
Asynchronous Command Delay	322
Output Formatting	
User-Friendly Command Output	323
Program-Friendly XML Output	
AVCLI Exceptions	
AVCLI Command Descriptions	
ad-disable	
ad-enable	335
ad-info	
ad-join	
ad-remove	338
alert-delete	
alert-info	
audit-export	341
audit-info	342
callhome-disable	
callhome-enable	
callhome-info	
datetime-config	
diagnostic-create	
diagnostic-delete	350
diagnostic-extract	351
diagnostic-fetch	352
diagnostic-info	
dialin-disable	
dialin-enable	
dialin-info	
disk-move-to-group	358
ealert-config	359
ealert-disable	
ealert-enable	
ealert-info	
help	
image-container-info	

image-container-resize	368
kit-add	369
kit-controlled-upgrade-continue	370
kit-controlled-upgrade-disable	371
kit-controlled-upgrade-enable	372
kit-controlled-upgrade-info	373
kit-delete	374
kit-info	375
kit-qualify	376
kit-upgrade	377
kit-upgrade-cancel	378
license-info	379
license-install	380
local-group-add	381
local-group-delete	382
local-group-edit	383
local-group-info	384
local-user-add	385
local-user-delete	.387
local-user-edit	388
local-user-info	390
localvm-clear-mtbf	391
media-create	.392
media-delete	393
media-eject	394
media-import	.395
media-info	.397
media-insert	398
network-change-mtu	399
network-change-role	400
network-info	401
node-add	403
node-cancel	.404
node-config-prp	405
node-delete	406

node-delete-prp	407
node-info	
node-reboot	
node-recover	410
node-shutdown	411
node-workoff	412
node-workon	413
ntp-config	
ntp-disable	415
ova-info	
ovf-info	417
owner-config	418
owner-info	419
pm-clear-mtbf	420
proxy-config	421
proxy-disable	
proxy-enable	423
proxy-info	
removable-disk-info	425
snmp-config	
snmp-disable	
snmp-info	
storage-group-create	430
storage-group-delete	431
storage-group-info	432
storage-info	
timezone-config	
timezone-info	435
unit-avoid-bad-node	436
unit-change-ip	437
unit-configure	439
unit-eula-accept	
unit-eula-reset	441
unit-info	
unit-shutdown	443

unit-shutdown-cancel	1
unit-shutdown-state	5
unit-synced	3
vm-attach-usb-storage	7
vm-ax-disable	9
vm-ax-enable450)
vm-boot-attributes	1
vm-cd-boot	2
vm-copy	3
vm-create	7
vm-create-from-snapshot	2
vm-delete	4
vm-device-config-info	5
vm-export	3
vm-import	7
vm-info)
vm-media-insert-disable471	1
vm-media-insert-enable	2
vm-network-disable	3
vm-network-enable	4
vm-poweroff	5
vm-poweron	3
vm-reprovision	7
vm-restore	1
vm-shutdown	4
vm-snapshot-create	5
vm-snapshot-create-disable	7
vm-snapshot-create-enable	3
vm-snapshot-delete	9
vm-snapshot-export)
vm-snapshot-info	2
vm-unlock	3
vm-usb-attach-disable	1
vm-usb-attach-enable	5
vm-volume-disable	3

vm-volume-enable	497
volume-info	
volume-resize	499
Chapter 11: System Reference Information	500
Compatible Guest Operating Systems	
Physical Machine System Requirements	502
Important Physical Machine and Virtual Machine Considerations	504
Virtual Machine Recommendations and Limits	505
Recommended Number of CPU Cores	505
Virtual Machine Limits	
Combined Virtual Machine Maximums	506
Important Considerations	
Creating a SplitSite Configuration	508
Creating the Configuration	512
A Typical everRun System	513
A SplitSite Configuration With a Quorum Server	513
SplitSite VLAN Requirements	514
From Initial Installation to Completing the SplitSite Configuration	515
Meeting Network Requirements	516
Locating and Creating the Quorum Server	518
Locating the Quorum Computer	518
Adding an Alternate Quorum Computer	519
Quorum Computer Requirements	
Downloading and Installing the Quorum Service Software	520
Completing the Configuration	
Configuring the Quorum Service Port	
Verifying the Quorum Service Port	
Configuring the Quorum Server Within the everRun Availability Console	522
Verify the Configuration and (Re-)Join VMs	
Understanding Quorum's Effect on System Behavior	
Example 1: A System Without a Quorum Server Experiences a Split-brain Condition	524
A Catastrophic Fault	524
Fault Handling	525
Recovery and Repair	525
Example 2: A SplitSite System With a Quorum Server Avoids a Split-brain Condition .	

A Catastrophic Fault	526
Fault Handling	527
Recovery and Repair	527
Example 2, Modified: The Quorum Server Is Unreachable During the Catastrophic Fault \ldots	528
Example 2, Modified: The Quorum Server Is Unreachable With No Catastrophic Fault \ldots	529
Recovering From a Power Failure	529
Accessing Knowledge Base Articles	529
Chapter 12: SNMP	532
Obtaining System Information with snmptable	532

Part 1: everRun User's Guide

The *everRun User's Guide* describes everRun systems, how to install them, and how to use them. For system descriptions including modes of operation and storage and network architecture, see:

• Introduction to everRun Systems

For planning and installation information, see:

Getting Started

The following topics describe how to administer everRun systems:

- Using the everRun Availability Console
- Upgrading everRun Software
- Managing Logical Disks
- Managing Physical Machines
- Managing Virtual Machines
- Maintaining Physical Machines

1

Chapter 1: Introduction to everRun Systems

See the following topics for an introduction to everRun systems:

- everRun System Overview
- Modes of Operation
- everRun Storage Architecture
- Network Architecture
- System Usage Restrictions

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

everRun System Description

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

Both PMs:

- Run the same host operating system (CentOS)
- Contain 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 and System Requirements Overview 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 Compatible Guest Operating Systems Network Architecture

Physical Machines and Virtual Machines

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

The everRun management software can create a guest VM from scratch. It can also import existing VMs from other environments and convert them into guest VMs. By creating an identical instance of the selected VM on a second host PM, the management software automatically provides high-availability (HA) or FT-class (based on the VM configuration) 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 guest VM—just as if the system administrator were loading them onto a physical server. If a fault or failure occurs in a disk or network device, the software automatically redirects I/O to the paired host PM for continuous operation. Though redundancy is lost until the failure is repaired, the VM continues to operate normally. The application continues to execute as if nothing had happened. The redundancy, fault detection, isolation, and management are completely transparent to the Windows or Linux environment and the application running within it. Repair of the PM is equally transparent and automatic. When a failed component on the PM is repaired, the software automatically incorporates the repaired components into the protected environment of the guest VM and restores redundancy transparently.

Related Topics

Using the everRun Availability Console The Physical Machines Page The Virtual Machines Page

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.

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.



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 for details.
- Proxy Configuration—Enables you to configure a proxy server for access to the Internet. See Configuring Internet Proxy Settings for details.

Lights-Out Management

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).
- If you install third-party management tools in the host operating system of a physical machine (PM) and you need to replace the PM in the future, remember to reinstall the tools on the replacement PM.

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

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

Related Topics

٥

Getting Started

System Reference Information

Modes of Operation

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

- High Availability Operation
- Fault Tolerant Operation

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.

High Availability Operation

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

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

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

Using the everRun Availability Console

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

Using the everRun Availability Console

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.

Meeting Network Requirements 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 and Locating and Creating the Quorum Server). 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

Creating a SplitSite Configuration

Network Architecture

Quorum Servers

A *quorum service* is a Windows operating system-based service deployed on a server distinct from the two servers (physical machines or PMs) of an everRun system. 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 the two PMs of an everRun system with 0, 1, or 2 quorum servers.

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

Quorum servers are particularly important in 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. See SplitSite Configurations for additional information.

If only two sites are available (thereby preventing the best practices configuration described above) and if one PM goes down and the surviving PM is unable to communicate with the quorum server (for example, because it is on the same site as the down PM), the VMs at the surviving site are automatically shut down to avoid running in split-brain.

Related Topics

Creating a SplitSite Configuration, which discusses quorum servers

Configuring Quorum Servers

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.

everRun systems support internal disks. 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.

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

- Logical Disks and Physical Disks
- Storage Groups
- Sizing Volume Containers
- Storage Requirements

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 ever-Run software detects new logical disks and logical disk failures. You manage logical disks using the everRun Availability Console. For information, see Managing Logical Disks.

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

The everRun Availability Console

Using the everRun Availability Console

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.

Related Topics

Storage Requirements

Creating a New Storage Group

The everRun Availability Console

Using the everRun Availability Console

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

0

Expanding a Volume Container on the everRun System

image-container-resize

Network Architecture

Ethernet networks provide pathways for communications between the two physical machines (PMs), or nodes, of the system. The main types of Ethernet networks are:

- One A-Link network must be a *private network* (priv0) that connects the two everRun PMs. For more information, see A-Link and Private Networks.
- Business networks allow your applications to connect to your existing network. One business network
 must be a management network (ibiz0, sometimes referred to as network0) that connects to the everRun Availability Console and is used by the quorum servers. For more information, see Business and
 Management Networks.

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.

A-Link and Private Networks

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

In addition to priv0, each system has A-Link networks to increase data-replication performance between the PMs. A-Link networks let you sync disks, shunt networks, migrate VMs, and sync fault-tolerant memory. By default, priv0 also performs the role of an A-Link network under the following conditions:

- If the priv0 speed is at least 10 Gb
- If the priv0 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 priv0 is not currently in use as an A-Link **and** is not the only remaining A-Link.

Priv0 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 priv0 at a later time.

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

Connect A-link networks between PMs either directly (that is, in the same manner that you connect priv0) or through a network switch.

Be sure to set up redundant A-Link networks.

The everRun installation software sets up priv0. 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.

Related Topics

Business and Management Networks

A-Link and Private Network Requirements

Network Architecture

Fixing a Network Connection

Business and Management Networks

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

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

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

You set up the management network when you install the everRun software. You can also set 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.

Related Topics

A-Link and Private Networks Business and Management Network Requirements Network Architecture Fixing a Network Connection

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 node and the stand-by node. The system's network segmentation logic detects a fault when this packet flow is interrupted while both sides still have an active network link. In this fault scenario, both nodes still have active network connections, so the fault lies in a switch that is external to the 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

System Usage Restrictions

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

- QEMU
- Accessing the Host Operating System

QEMU

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 you install the everRun software, you can access the host operating system (CentOS) locally at the PM's physical console, or you can access it remotely by using a secure shell (SSH) client.

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

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

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

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

П

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

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

To locate the IP address of each PM in the everRun Availability Console

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

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

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

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

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

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

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

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

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

Chapter 2: Getting Started

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

- Planning
- Software Installation
- Post-Installation Tasks

Planning

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

- System Requirements Overview
- Storage Requirements
- Memory Requirements
- General Network Requirements and Configurations
- Business and Management Network Requirements
- A-Link and Private Network Requirements
- everRun Availability Console Requirements
- Compatible Internet Browsers
- Power Requirements and Considerations
- Creating a SplitSite Configuration (if applicable to your configuration)

After you have planned the system configuration, continue with Software Installation.
System Requirements Overview

An everRun system requires two x86-64 host servers (physical machines (PMs) or nodes) 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 requirements are summarized below.

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

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.

A second computer with identical processors is required for use as a redundant server for guest Virtual Machines (VMs), which 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

Internal 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 the guest VMs may need. If planning to run multiple VMs, consider adding pairs of A-Links, up to the supported total of four pairs.

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

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

IP Addresses

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

Ports

everRun systems use port 443 in the local firewall for HTTPS communications, port 22 for ssh, and 5900-59*nn* for each active VNC associated with each VM. Firewalls must allow traffic through the appropriate ports. Firewalls must permit VMs to contact quorum service computers using UDP port 4557. For additional information on TCP and UDP ports, access the Knowledge Base to search for the article *TCP and UDP ports used by everRun* 7 (KB-2123). See Accessing Knowledge Base Articles.

Related Topics

Physical Machine System Requirements Important Physical Machine and Virtual Machine Considerations Virtual Machine Recommendations and Limits Planning Virtual Machine Resources Configuring IP Settings

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.
 - 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. The everRun software does not monitor the state of physical disks in a RAID set.

Disk drives supported include 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. everRun systems support 4K sector size disks in native mode. 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.

In addition, note the following restrictions of the Initial Storage Group:

- If you add a second logical disk to the default **Initial Storage Group**, its size must be greater than 32.2GB.
- everRun software prevents all non-volatile memory express (NVMe) devices from being included in the Initial Storage Group because not all BIOS and UEFI systems allow NVMe devices to be bootable devices.

When planning your system configuration, confirm that your storage configuration meets these requirements and then return to Site and System Preparation.

Related Topics

everRun Storage Architecture Managing Logical Disks The Storage Groups Page

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

General Network Requirements and Configurations

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

Note: SplitSite networks have some additional and different network requirements and recommendations. See Creating a SplitSite Configuration in addition to the information below.

Requirements

٥

Before you install the 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
- Business and Management Network Requirements

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.

П

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.
- Virtual machines (VMs) can use IPv4, IPv6, and other Ethernet protocols.
- All business networks can be used for IPv6 host access if your site has SLAAC or DHCPv6 enabled.
- To reach the everRun Availability Console, use ibiz0, which is the IPv4 address that migrates to the primary management physical machine (PM). Each PM also has its own ibiz0 IPv4 address on the management network.
- Each PM requires at least one business network (specifically, the management network), 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 PM.
- The switch ports connected to business networks must allow layer2 multicasts (address: 01:E0:09:05:00:02) with ethertype: 0x8807.
- If you configure RHEL or CentOS guests to have multiple NICs on same subnet, you may experience guest network connectivity issues due to asymmetric routing. To avoid this problem, modify the /etc/sysctl.conf file on the guest Virtual Machine (VM) to contain the following lines, save the file, and reboot the VM.
 - net.ipv4.conf.default.rp filter = 2
 - net.ipv4.conf.all.rp filter = 2
- Do not issue the ifdown command from a PM's host OS to temporarily bring down a VM's business network connection (ibiz x). Doing so will disconnect the physical interface from its bridge and cause the VM to become unreachable over the network. Instead, use the ifconfig down command.
- The switches connected to business networks must not enable any MAC address security features that would disable the movement of a MAC address from one business link to the matching business link on the other PM.
- For optimal failover response, configure any switches connected to your system to have MAC aging timeout values of less than one second.

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

Related Topics

Network Architecture

Business and Management Networks

network-info and network-change-mtu

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 physical machine (PM), without any protocol filtering.
- Ethernet packets sent between two PMs of a system must not be obstructed or rate-limited. Ensure that they are not routed or switched by any L3 network infrastructure.
- 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. A-link networks can use 25 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 virtual machine (VM) to a minimum of <u>one</u> or a maximum of <u>two</u> A-Link networks. However, each A-Link network can have multiple VMs assigned to it.

You may be able to improve VM performance and reduce host processing overhead by enabling jumbo frames on A-Link networks. To do so, change their Ethernet frame MTU size from the default 1500 bytes to 9000 bytes. For instructions, access the Knowledge Base to search for the article *Optimizing Performance of ever-Run A-Link Networks* (KB-4262). See Accessing Knowledge Base Articles.

Related Topics

A-Link and Private Networks

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 management network.
- Use a supported browser. See Compatible Internet Browsers.

For more information, see Using the everRun Availability Console.

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.

Power Requirements and Considerations

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

See Connecting Power 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.
- 2. Connect power to your system. See Connecting Power.
- 3. Install the everRun software. See Installing everRun Software.

When the installation is complete, see Post-Installation Tasks.

Related Topics

Upgrading everRun Software

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.
- The storage configuration meets all of the requirements described in Storage Requirements.
- 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.
- Determine the best configuration for your network. See General Network Requirements and Configurations.
- Use either an internal DVD drive or bootable USB media (see Creating Bootable USB Media) for the installation.

After confirming that your site and system meet the preceding requirements, return to Software Installation.

Connecting Power

After connecting power, return to Installing everRun Software.

UPS (Optional)

П

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

Single UPS:



Dual UPS:



Related Topics

Power Requirements and Considerations

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

- From any computer connected to the Internet, go to the Downloads page at https://www.stratus.com/services-support/downloads/?tab=everrun.
- 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). Then, perform the next step in Installing everRun Software.
- If you are not creating bootable media, perform the next step in Installing everRun Software.

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 unmount 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-7.8.0.0-192.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.



L

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/.

- Download and save the everRun software ISO image (see Obtaining everRun Software) to a Windows-based system, if you have not already done so.
- Ensure the integrity of the ISO image using a tool for verifying MD5 checksums on Windowsbased 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.
- 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

Option	Value(s)
	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 ever-Run.

When the USB device is ready to be used to install everRun software, perform the next step in Installing ever-Run Software.

Related Topics

Obtaining everRun Software

Software Installation

ñ

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).

After you modify the settings, save them and perform the next step in the installation procedure (either Installing Software on the First PM or Installing Software on the Second PM).

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 <u>required</u>.

First Boot Device	 Controls which device boots the operating system. Set the first boot device to the appropriate value for the boot device you are using: Optical Drive for a DVD drive The appropriate value for a flash drive (for example, USB Storage or USB Device)
Virtualization Technology	Allows the processor to use virtualization technology. Set this to Enabled.
Execute-Dis- able Bit Cap- ability	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.

The
ever-

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
- 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).
 - b. Connect Ethernet cables for the networks you are configuring. See Connecting Ethernet Cables.
- 3. Perform the installation on the first PM. See Installing Software on the First PM.
- After you have finished installing the software on the first PM, perform the installation on the second PM. See Installing Software on the Second PM.
- The software installation is complete. Now, perform the required post-installation configuration steps. See Post-Installation Tasks.

Connecting Ethernet Cables

٥

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

Note: To install additional networks <u>after</u> you have completed the software installation, see Connecting Additional Networks.

On each physical machine (PM), assign one network port as the private network (**priv0**), and assign another network port as the management network (**ibiz0**, sometimes referred to as network0). 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.

1 I

П

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 <u>one</u> A-Link network in addition to the private network. See A-Link and Private Network Requirements.

• 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.

Related Topics

Software Installation A-Link and Private Network Requirements Business and Management Network Requirements everRun Availability Console Requirements

Installation Options

When you begin the installation and you have selected a keyboard map, a 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 per- form 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 install- ation medium the first time you use it; note, though, that verification adds five or so minutes to the installation.) See Installing Software on the First PM.
Verify the installation medium and then recover a physical machine.	Verify medium and Recover Physical Machine	Verifies the installation medium and then recovers a phys- ical machine. See Recovering a Failed Physical Machine.
Verify the installation medium and then replace a physical machine.	Verify medium and Replace Physical Machine	Verifies the installation medium and then replaces a phys- ical machine. See Replacing Physical Machines, Mother- boards, NICs, or RAID Controllers.

Task	Option	Description
Perform initial install- ation 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 sys- tem. See Installing Software on the First PM.
Recover failing PM	Recover PM, Join system: Preserving data	Preserves all data but re-creates the /boot and root file systems, reinstalls the CentOS and everRun soft- ware, and attempts to connect to an existing system. (This option is the default.) See Recovering a Failed Phys- ical Machine.
Perform initial install- ation on the second PM; replace a PM	Replace PM, Join system: Initialize data	Deletes all partitions on all connected disks, installs CentOS and everRun software, and attempts to connect to an existing system. See Installing Software on the Second PM and Replacing Physical Machines, Mother- boards, NICs, or RAID Controllers.
Boot to rescue mode (UEFI firmware install- ation only)	Rescue the installed system	Boots to rescue mode.

Installing Software on the First PM

П

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.
- 3. When the installation software loads, the Welcome to everRun release_number window appears with a list of keyboard map values and other options. Immediately below the options list, text describes the keys you use to select an option, and Help text (BIOS mode) appears below the key descriptions. Use arrow keys to select one of the following options:
 - The country keyboard map for the installation—Use the arrow keys to select one of the appropriate values, and then press Enter:

Note: If you need to set a keyboard map for a different country, see Mapping Your Keyboard.

Country	BIOS Mode Value	UEFI Mode Value	keyboard map = (BIOS) keymap = (UEFI)
China	China	中国	NONE
Germany	Germany	Deutschland	de
Japan	Japan	日本	jp106
United States (default)	U.S.A.	U.S.A.	us

- Troubleshooting Utilities (BIOS firmware installation, only)—When you select this option, the Troubleshooting everRun release_number window appears. Use the arrow keys to select one of the following tasks:
 - **Rescue the installed system**–Boots to rescue mode.
 - Boot from the local disk drive (the default)–Boots from a local disk drive.
 - **Memory test**–Performs a memory test.
 - Return to the Installer main menu–Returns to the Welcome window.

Help text appears at the bottom of the window for some troubleshooting selections. The system performs the task you select; the remaining steps in this installation topic are not relevant.

If you select a country keyboard map option, continue with the next step.

4. The Install or Recover (keymap) everRun release_number window displays the list of options shown in Installation Options. Immediately below the list, text describes the keys you use to select an option. In BIOS mode, Help text appears below the key descriptions.

In this window, choose one of the following methods to perform the initial installation:

- 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

ň

Note: At any point when using the user interface, you can press the **Tab** key (BIOS mode) or the **E** key (UEFI mode) to display and edit the command line (Method 2: Installing via the Command Line).

i. In the **Install or Recover** (*keymap*) everRun *release_number* window, use the arrow keys or the highlighted letters to select an installation option.

Stratus recommends that, for an initial installation, you select Verify medium and Install ever-Run, which verifies the installation medium before installing the software, adding about five minutes to the installation process. If you select verification, the system displays checking: nnn.n% and various other messages. If successful, installation continues. If verification fails, the installation stops. After you have verified the medium once, you do not need to verify it again. If you do not want to verify the medium, select Install everRun, Create a new system.

After selecting an installation option, press Enter. The installation proceeds:

- BIOS mode—Many messages appear on the screen and sometimes a short delay occurs.
- UEFI mode—The screen is blank for several seconds, and then messages appear on the screen.

Note:

ñ

If any disk contains previously installed data, various messages appear, including the following, and the system reboots (if no disk contains previously installed data, installation continues):

DISKS WERE WIPED. REBOOTING TO RESTART THE INSTALLER.

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 4, above).

ii. The Select interface for private Physical Machine connection dialog box appears, enabling you to select the physical interface for the private network (priv0). The first embedded port, em1, is selected by default. Use the arrow keys to navigate to another port, if necessary. Press the space bar to select the interface, and then press the Tab key to navigate to OK, which saves the selection and continues the installation.

Notes:

П

- 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.
- iii. The Select interface for managing the system (ibiz0) dialog box appears with a list of available physical interfaces for the management network. Use the arrow keys to navigate to the second embedded port, em2 (if it is not already selected). Press the space bar to select it, and then press the Tab key to navigate to OK, which saves the selection and continues the installation (or use arrow keys to navigate to Back, to return to the previous 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.

- iv. The Select the method to configure ibiz0 dialog box appears, enabling you to set the management network for node0 as either a dynamic or static IP configuration. Typically, you configure ibiz0 as a static IP configuration. Use the arrow keys or the Tab key to navigate to one of the following options:
 - Automatic configuration via DHCP–Select this option to configure ibiz0 as a dynamic IP configuration. Press the space bar to select the option and then press F12 to save your selection.
 - Manual configuration (Static Address)—Select this option to configure ibiz0 as a static IP configuration. Press the space bar to select the option and then press F12 to save your selection. The Configure em2 dialog box appears. See your network administrator for the information you enter in this dialog box. Enter the following information:
 - IPv4 address
 - Netmask
 - Default gateway address
 - Domain name server address

Error messages appear in the window if values are incorrect.

After you enter the information, press F12.

The **Confirm configuration choices** dialog box appears. Use the arrow keys or the **Tab** key to navigate to **OK** (to use the displayed values) or to **Back** (to return to the previous dialog box, to change your selections).

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

Method 2: Installing via the Command Line

Note: To return to the **Install or Recover** window from the command line, press the escape (Esc) key.

- i. Press the **Tab** key (BIOS mode) or the **E** key (UEFI mode) to display and edit the command line.
- ii. Set the value for the private network (**priv0**) by typing one of the following values.
 - To use the first embedded interface:

priv0=em1

• To automatically select the default interface:

priv0=auto

• To use the interface with a MAC address:

priv0=AA-BB-CC-DD-EE-FF or priv0=AABBCCDDEEFF

- iii. Set the value for the management network (ibiz0) by typing one of the following values.
 - 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 typing values on the command line, press Enter.
- v. If any disk contains previously installed data, various messages appear, including the following, and the system reboots (if no disk contains previously installed data, installation continues with the next step):

DISKS WERE WIPED. REBOOTING TO RESTART THE INSTALLER.

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 4, above).

- 5. 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 <u>dynamically</u>, record its IP address as described in Recording the Management IP Address.
- 6. Perform the next step in Installing everRun Software.

Mapping Your Keyboard

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

Supported keyboard layouts include:

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

To configure your keyboard layout during installation:

- 1. As the first PM boots, select a keyboard map on the **Welcome** menu, press **Enter**, and in the next screen, select **Install**, **Recover**, or **Replace**.
- 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 Swiss German keyboard layout:

inst.keymap=sg

- On legacy BIOS systems, press Enter to continue the boot sequence. On UEFI systems, press Ctrlx.
- 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.
- 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

Post-Installation Tasks

Recording the Management IP Address

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

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

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 to continue installation.

Related Topics

П

Business and Management Network Requirements

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), using the user interface.

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.
- 3. When the installation software loads, the **Welcome** screen appears with a list of keyboard map values and other options. Select the country keyboard map or other option you require. This topic describes how to perform an installation after selecting a country keyboard map. To perform an installation with a flash drive, see Installing Software on the First PM.
- 4. The Install or Recover ... screen appears and displays the options shown in Installation Options. 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.
- 5. 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.

6. 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.

 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:

Ö

П

- 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.

9. 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.

- 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 redisplays until you enter valid information.

- 11. 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 <u>dynamically</u>, record its IP address as described in Recording the Management IP Address
- 12. Perform the next step in Installing everRun Software.

Post-Installation Tasks

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

- Obtaining System IP Information
- Logging On to the everRun Availability Console for the First Time
- Configuring Required System Preferences:
 - Configuring Date and Time
 - Configuring Remote Support Settings
 - Configuring Quorum Servers
 - Specifying Owner Information
- Configuring Active Directory
- Managing Local User Accounts

- Resolving Outstanding Alerts on the Dashboard
- Connecting Additional Networks

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). 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), 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

Post-Installation Tasks

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 provide network information. You can also acquire a permanent license now, though you can do so later. When a system is first installed, it has a temporary license that expires within 30 days.

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

- The node0 (primary) IP address—You record this address during installation. See Recording the Management IP Address.
- The system IP address—The network administrator provides this information. See Obtaining System IP Information.
- The temporary license file (*site-id_L.KEY*) that you received from Stratus when you purchased the everRun software—Your company typically receives this file in email, or you can download it from the Stratus Customer Service Portal at https://support.stratus.com.

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 log-on page of the everRun Availability Console appears.

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

The Stratus everRun END USER LICENSE AGREEMENT (EULA) appears.

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

The INITIAL CONFIGURATION page appears under Config.

- 4. Under NOTIFICATIONS, the box for Enable Support Notifications is checked, by default. If you do not want the 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).
- Under SYSTEM IP, for Static System IP, enter the static system IP address that you obtained from your network administrator. (The system IP address is sometimes referred to as the cluster IP address.)
- 6. Also under **SYSTEM IP**, select **DHCP** (the default) or **Static**. For **DHCP**, you do not need to provide additional information.

If you select **Static**, the node0 static IP address that you entered during deployment appears. Provide the following values:

- Primary and secondary DNS
- NetMask
- Gateway address for node0
- IP address for node1
- Gateway address for node1

Confirm that the IP address for the management network (ibiz0) is correct.

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

LICENSE INFORMATION window appears.

 You can upload a new license key now or later from the Product License page. To do so now, click Choose File in the LICENSE INFORMATION window under Upload License Key. Navigate to the license .KEY file that you received from Stratus. Select the license file and click Upload.

To do so later, click **Finish**. See Managing the Product License for information on uploading a license.

8. Click Finish.

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

Perform additional tasks in Post-Installation Tasks, if necessary.

Related Topics

Software Installation

The everRun Availability Console

Using the everRun Availability Console

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

- 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

A-Link and Private Network Requirements

Business and Management Network Requirements

General Network Requirements and Configurations

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.

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

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

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.

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

- Read system alerts from the Dashboard. See The Dashboard Page.
- View VM, CPU, memory, and storage statistics, and reboot or shutdown the system from the System page. See The System Page.
- Set preferences for the system, notifications (e-Alerts and SNMP configuration), and remote support (notification and access); and access administrative tools that enable you to 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, active directory, etc. See The Preferences Page.
- View alerts and audit logs. See The Alerts History Page, The Audit Logs Page, and The Support Logs Page.
- Monitor, manage, and maintain resources:
 - PM status, storage (including disks), network, VMs, and USB devices: see The Physical Machines Page.
 - VM status and management tasks such as creating, importing/restoring, managing, and maintaining VMs: see The Virtual Machines Page.
 - Snapshot status and management tasks such as exporting and deleting snapshots: see The Snapshots Page.
 - Volumes, including their state, name, data synchronization status, size, storage group, state, and other information: see The Volumes Page.

- Storage groups, including name, size used, size, and number of volumes: see The Storage Groups Page.
- Networks, including state, link condition, name, internal name, type (for example, A-Link), VMs, speed, MAC address, and network bandwidth: see The Networks Page.
- Virtual CDs, including their storage group, state, name, size, and whether or not the VCD can be removed: see The Virtual CDs Page.
- Monitor and manage upgrade kits, and users and groups in the LIBRARY. See The Upgrade Kits Page and The Users & Groups Page.

Related Topics

Logging On to the everRun Availability Console for the First Time

Logging On to the everRun Availability Console

Using the everRun Availability Console

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 view alerts and logs, and perform other administrative tasks.



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

The everRun Availability Console

Using the everRun Availability Console

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, ⁽¹⁾) in the ever-Run 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).

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 The System Page

The Virtual Machines Page

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

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 statistics and 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
- Shutting Down the System

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

To manage everRun system resources, see Configuring System Resources.

Viewing statistics

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

- Virtual Machines-A table displays the State, Activity, and Name of each VM.
- *system name*–Circle graphs indicate the system's CPU allocation, memory allocation, disk (R/W), and network utilization.

- Node0 and Node1–Circle graphs indicate each node's CPU utilization, memory utilization, disk utilization, and network utilization. For disk utilization and network utilization, you can select the logical disk or the network whose statistics you want to display.
- PM Detail and VM Detail—Line graphs display the percentages of total CPU capacity, total memory capacity (for PMs only), disk I/O (in bytes/s), and network I/O (in bits/s). You can select the time span for the statistics in the range of the last four hours to last year. You can also select to display live statistics.

At the far right of each heading, you can select the entity whose statistics you want to display. For example, under **PM Detail**, to the far right of **Percentage of Total CPU Capacity (%)**, you can select **node0** or **node1** in the **Physical Machines** drop-down box.

Click the arrow to the left of each heading to expand or collapse the display.

Related Topics

٥

Using the everRun Availability Console

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).

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

I

The everRun Availability Console

The System Page

Using the everRun Availability Console

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:

I

- 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. Select System in the left-hand navigation panel.
- 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 <u>Maintenance Mode</u>). 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 shutdown completely, a VM may not be shutting down properly. Do one of the following to shut down the VM:

- Use the VM console or a remote desktop application to log on to the VM. Use operating system commands to shut down the VM.
- Log on to the 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

The everRun Availability Console

The System Page

Using the everRun Availability Console

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.

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.
Product License	Allows you to view and manage the everRun product license. See Managing the Product License.

The following table lists and describes the preferences.

Preference	Description
IP Configuration	Allows you to view and specify the Internet Protocol (IP) address and net- work settings for the system. See Configuring IP Settings.
Quorum Servers	Allows you to view existing and new Quorum servers. Quorum servers provide data integrity assurances and automatic restart capabilities for spe- cific failures in the everRun environment. See Quorum Servers and Con- figuring Quorum Servers.
Date & Time	Allows you to view the system time, specify values for Network Time Pro- tocol (NTP) (recommended), or to manually set the time and date on the sys- tem. See Configuring Date and Time.
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.
Administrative Tools	
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.
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.
Secure Connection	Allows you to enable only HTTPS connections to the system. See Con- figuring Secure Connections.
Host Inactivity Logout	Allows you to disable the host inactivity logout or to change the timeout. See Configuring the Host Inactivity Logout.

Preference	Description	
Snapshot Configuration	Allows you to disable the creation of snapshots. See Disabling and Enabling Snapshots.	
VM Device Configuration	Allows you to disable or enable insertion of virtual CDs (VCDs) in all VMs or attachment of USB devices to all VMs. See Configuring VM Devices.	
IPtables Security	Allows you to manage IP packet filtering using the administrative tool IPt- ables. See Managing IPtables.	
Notification		
e-Alerts	Allows you to enable email alerts (e-Alerts) for system administrators. See Configuring e-Alerts.	
SNMP Configuration	Allows you to enable Simple Network Management Protocol (SNMP) requests and traps for remote system monitoring. See Configuring SNMP Settings.	
Remote Support		
Support Configuration	Allows you to configure remote access and notifications. Remote access enables your authorized Stratus service representative to log on to the sys- tem 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.	
Proxy Configuration	Allows you to configure proxy settings for the everRun system if your organ- ization requires a proxy server to access the Internet and you have a service agreement with Stratus or another authorized everRun service rep- resentative. The everRun software uses proxy server information for support notification messaging and remote support access features. See Configuring Internet Proxy Settings.	

Related Topics

The everRun Availability Console

Using the everRun Availability Console

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

The everRun Availability Console

Managing the Product License

Manage the product license for the system by:

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

When you purchase a 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 for the first time.

Once a system has a permanent license, it checks with the license server for updates every 24 hours, if the system has an Internet connection. If a system does not have Internet access, you can still update the license and check its status. To do so, you need to move a file between the location of the everRun

Availability Console (which does not have Internet access) and a location with Internet access. Two methods are as follows, though other methods are possible:

- A USB flash drive—You move a USB flash drive between a management PC (which can connect to the system) and a computer with Internet access.
- A mobile device such as a laptop or smart phone—You move a mobile device between a location where you can log in to the everRun Availability Console and a location with Internet access.

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

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

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

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

- a. On the Preferences page, click Product License.
- b. Click the License Check and Activation bar.
- c. For License Check, click Check License Now.
- 2. The console displays the status of the license (date format varies, based on location):

211772	License is activated and does not
SIAIUS	expire.
TYPE OF LICENSE	Enterprise Edition (volume)
EXPIRATION	day, month dd, 20yy, time
LAST CHECK	day, month dd, 20yy, time
ASSETID	asset_ID
FT Enabled	Yes_or_No
Split Site Allowed	Yes_or_No

To update a new license manually

After you have saved a license.key file to a computer, use this procedure to upload the license.key file to the everRun system.

On a system with Internet access

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

- 2. On the **Preferences** page, click **Product License**.
- 3. Click the **New License** bar to display its options.
- 4. Under Upload New License Key, click Choose File, 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.

On a system without Internet access

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

On a computer or mobile device with access to the everRun Availability Console

1. If using a management PC, insert a USB flash drive into a USB port.

If using a mobile device, ensure that it has access to the everRun Availability Console.

- 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.
- 6. Under Step 1, Download Activated License Key, click Save activation URL file and navigate to the USB flash drive or mobile device, to save the file. Alternatively, click the Activated License link and select your browser's option to copy the link (for example, Copy Link Location or Copy Link Address). Open a text editor (notepad.exe), paste the copied URL into it, and save the contents of the editor to a text file on the USB flash drive or mobile device. If using a USB flash drive, remove it. Go to a location with Internet access.

In a location with Internet access

- 1. If using a USB flash drive, insert it into a USB port of the computer with Internet access.
- 2. Navigate to the file you saved, and click the file name. 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. If using a USB flash drive, copy the new license .key file to it, and then remove the USB flash drive.
- 4. Return to the location with access to the console.

On a computer or mobile device with access to the everRun Availability Console

1. If using a USB flash drive, insert it into a USB port on the management PC.

If using a mobile device, ensure that it has access to the everRun Availability Console.

- 2. In the console, click **Preferences** in the left-hand navigation panel.
- 3. On the **Preferences** page, click **Product License**.
- 4. Click the License Check and Activation bar.
- 5. Under Step 2, **Upload Activated License Key**, click **Choose File** and navigate to the location where you saved the file.
- 6. Select the file, click **Open**, and then click **Upload** to upload the file to the system.

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

To view the license activation 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, the 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, 2). If necessary, your authorized Stratus service representative can change the limit.

9.0: ALAS_UNKNOWN

Unknown error.

Related Topics

The everRun Availability Console

The Preferences Page

Using the everRun Availability Console

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 the everRun software, you configure three IP addresses: one for the system and one for each node (node0 and node1). 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 ever-Run system.

Warnings:

- 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.
- You must use the everRun Availability Console to change IP addresses. Do not use Linux tools.

Notes:

Ö

- 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 different 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. In the procedure that the article describes, you have the option of using the Save and Shutdown button in the IP Configuration section of the Preferences page.
- 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.

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

Related Topics

Software Installation

Obtaining System IP Information

Logging On to the everRun Availability Console for the First Time

The Preferences Page

The everRun Availability Console

Using the everRun Availability Console

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 and Creating a SplitSite Configuration (which discusses quorum servers).

Notes:

Ö

- For a VM to recognize quorum server configuration changes, you must reboot the VM by shutting it down and then restarting it. See Shutting Down a Virtual Machine and Starting a Virtual Machine.
- 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

٥

The everRun Availability Console

The Preferences Page

Using the everRun Availability Console

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

П

Ö

- 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

The Preferences Page

Using the everRun Availability Console

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

- 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

The Preferences Page

Using the everRun Availability Console

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).

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 the Enable Active Directory button.
- 4. Next to Active Directory Domain, type the name of the domain to use.
- 5. Click one of the following to prevent or allow automatic assignment of 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. Type the **Username** and **Password** of an Active Directory Administrator in order to add this everRun system to the domain.
- 7. Click Add.
- Assign administrative privileges to domain users on the Users & Groups page, as described in Managing Domain User Accounts.

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 Managing Domain User Accounts Managing Local User Accounts The Preferences Page The everRun Availability Console

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 (⁴¹) appears in the masthead with the message **VMs are not load balanced** and a link to Load Balancing. Click the link to rebalance the load.

Related Topics

Managing Virtual Machines

The Preferences Page

The everRun Availability Console

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:

П

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,

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

To enable only HTTPS connections

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

The system automatically logs you out of the 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.

To install a custom certificate

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

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

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

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

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

• Recover the custom certificate to the previously used one:

```
certificate installer recover -p
```

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

certificate installer recover -b

· List information about the currently used certificate:

```
certificate_installer list -c
```

• List information about the previously used certificate:

certificate_installer list -p

If you want more information about installing a custom certificate, access the Knowledge Base to search for the article *Adding Certificates to ca-bundle.crt in everRun Enterprise* (KB-4284). See Accessing Knowledge Base Articles.

The certificate_installer script

Usage

certificate_installer [command_options] [script_options]

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

Script Options

-v,[no_]verbose	In verbose mode, the script displays all inform- ation.
-1,log= <i>log_file</i>	Prints logs to the file <i>log_file</i> instead of to STDOUT.

Related Topics

The everRun Availability Console

The Preferences Page

Using the everRun Availability Console

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

The everRun Availability Console

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

Managing Snapshots

Using the everRun Availability Console

Configuring VM Devices

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

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

To disable insertion or attachment of VM devices

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

To enable insertion or attachment of VM devices

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

Related Topics

Inserting a Virtual CD

Attaching a USB Device to a Virtual Machine

The Preferences Page

Using the everRun Availability Console

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 **IPt-ables 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 the host operating system on each physical machine (PM), to both IPv4 and IPv6 packets, and the rules remain persistent after rebooting. When you insert 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:

П

- For information on the ports that everRun software uses, see System Requirements Overview.
- 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.

To manage IPtables, first, enable IPtables security, if you have not already done so.

To enable IPtables security

- 1. Click **Preferences** in the left-hand navigation panel, to open the **Preference** page.
- 2. On the **Preferences** page, click **IPtables Security**.
- 3. Activate the checkbox next to Enable Port Management.

The **Enable Port Management** window becomes gray for a few minutes. When the window is active again, **Enable Port Management** is selected

Continue, as appropriate, by inserting a new rule, removing a rule, loading default settings, importing rules, or exporting rules.

To insert 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**.

Ensure that Enable Port Management is selected.

- 3. Click the Insert New Rule button to open the Insert New Rule pop-up window.
- 4. 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.
- Port Number (starting)—For the first port of the range, enter a number 0 to 65535 that is less than or equal to Port Number (ending).
- Port Number (ending)—For the last port of the range, enter a number 0 to 65535 that is greater than or equal to Port Number (starting).

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

5. 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**.

Ensure that Enable IPtables Security is selected.

- 3. Select the rule that you want to remove.
- 4. Click **Remove** (in the right-most column), for the rule you selected.
- 5. 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**.

Ensure that Enable Port Management is selected.

3. 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. If you click **OK**, the **Enable Port Management** window becomes gray for a few minutes and the *Load-ing default settings....* message appears.

Related Topics

The Preferences Page

The everRun Availability Console

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.)
 - **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. Then, type the Username and Password for the SMTP account.
- List of Recipients (required)–Enter email addresses for all e-Alert recipients.
- 4. Click **Save** (or click **Reset** to restore the previously-saved values).

٥

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

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

Related Topics

The everRun Availability Console

The Preferences Page

Using the everRun Availability Console

Configuring SNMP Settings

Configure Simple Network Management Protocol (SNMP) settings for the 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 system to retrieve the values of objects listed in the Management
 Information Bases (MIBs) supported by the everRun software. These MIBs include a system-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 https://www.stratus.com/services-support/downloads/?tab=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).

Follow the appropriate procedure to enable SNMP requests or traps.

To enable SNMP requests

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

To customize the list

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

To enable SNMP requests

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

The selections in Step 4 appear as well as a read-only display of the snmpd.conf file.

4. Select one of the following:

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

Unrestricted–Allows public access.

Customized (available when snmpd.conf has been edited)-Allows customized access.

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

To enable SNMP traps

П

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

Related Topics

SNMP

The everRun Availability Console

The Preferences Page

Using the everRun Availability Console

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

- 1. Click **Preferences** in the left-hand navigation panel, to open the **Preferences** page.
- 2. Under Remote Support, click Support Configuration.
- 3. Modify the settings, as appropriate for your system:

- Enable Remote Support Access allows your authorized Stratus service representative to remotely connect to the 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.x* (KB-2123). See Accessing Knowledge Base Articles.)
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); and SNMP sends traps to recipients of SNMP traps, if configured (see Configuring SNMP Settings). Watch the Alerts History log (see The Alerts History Page) for delivery status. A subsequent alert will be generated if the support notification fails.

Related Topics

The everRun Availability Console

The Preferences Page

Using the everRun Availability Console

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).

- 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).

The everRun Availability Console

The Preferences Page

Using the everRun Availability Console

The Alerts History Page

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

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

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
 - Configuring e-Alerts
 - Configuring SNMP Settings

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

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

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

Related Topics

Ö

The everRun Availability Console

Using the everRun Availability Console

The Audit Logs Page

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

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

- **Time**—The date and time of the action.
- Username-The name of the user that initiated the action.
- Originating Host-The IP address of the host on which the 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).

Related Topics

The everRun Availability Console

Using the everRun Availability Console

The Support Logs Page

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

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

The everRun Availability Console

Using the everRun Availability Console

The Preferences Page

Creating a Diagnostic File

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

Note: 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

ñ

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

Related Topics

The everRun Availability Console

The Preferences Page

Using the everRun Availability Console

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 system. (To create a diagnostic file, see Creating a Diagnostic File.)

To upload a diagnostic file to Customer Support

- 1. Click Support Logs in the left-hand navigation panel, to open the Support Logs page.
- 2. 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. If the upload succeeds, a message appears, confirming that the diagnostic file was uploaded successfully.
 - 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 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 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 https://www.stratus.com/services-support/customer-support/?tab=everrun.

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

Related Topics

The everRun Availability Console The Preferences Page Using the everRun Availability Console

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

Related Topics

The everRun Availability Console

The Preferences Page

Using the everRun Availability Console

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. To display information about a node, you can use the snmptable command; see Obtaining System Information with snmptable.

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

- Action buttons: Various action buttons appear, with inactive buttons grayed out, depending upon the state of the selected node. 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). To learn about additional PM actions available in maintenance mode, see Physical Machine Actions or the help topic for the task you want to complete.
- Detailed information: To view detailed information or statistics about the selected node, click one of the following tabs:

- Summary (in the initial display), which displays information about the node, such as (if applicable) the manufacturer, the model, serial number, overall state, activity, and configuration (memory and logical disks) for the selected node.
- Description, which displays a window where you can enter information about the node.
- Storage, which displays the state, logical ID, sector size, size, and size used of storage. It
 also displays the storage group and the current action (if any).
- **Network**, which displays the state, name, speed, and MAC address of networks.
- Virtual Machines, which displays the state, activity, and name of virtual machines.
- USB Devices, which lists any USB devices inserted in the node.

L

The everRun Availability Console

Using the everRun Availability Console

Physical Machine Actions

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

Caution: Use the **Physical Machines** page of the everRun Availability Console when you perform maintenance on a PM. Do not use controls on the computer (for example, the power switch on the PC), because the everRun Availability Console protects the everRun system from most actions that are potentially disruptive.

Commands	Description	
Work On	Enters a PM into maintenance mode. VMs running on this PM migrate to the other PM, if it is in service. (Otherwise, you are asked to re-confirm the request and then shut down VMs.) When VMs are migrated or shut down, the PM displays running (in Maintenance) . See Maintenance Mode.	
The following actions are available after clicking the Work On button, when the PM has entered main-		

Commands	Description		
Finalize	Removes a PM from the state running (in Maintenance) . See Maintenance Mode.		
Shutdown	Shuts down a PM. The PM transitions to off (in Maintenance) . See Shutting Down a Physical Machine.		
Reboot	Reboots the PM. The PM transitions to preparing for reboot (in Maintenance) . See Rebooting a Physical Machine.		
Remove	Causes the everRun software to delete the PM from the everRun system's data- base, so that you can replace the PM or one of its components. See Replacing Physical Machines, Motherboards, NICs, or RAID Controllers.		
The following action may be available when a PM has failed or when the everRun software has removed a PM from service and powered it off, due to an excessive failure rate.			
Recover	Recovers a failed PM. In some cases, the everRun Availability Console disp the state of a failed PM as Unreachable (Syncing/Evacuating) . See Rec Recover ering a Failed Physical Machine.		

The everRun Availability Console

Using the everRun Availability Console

The Physical Machines Page

Physical Machine States and Activities

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

State	Activity	Available Commands	Description	
*	Running	Work On	PM is running normally.	
Q	Evacuating	Finalize	Virtual machines are migrating from this PM to its partner.	
1	Running	Work On	PM is predicted to fail.	
×	Running	Work On	PM failed.	
×	× Powered Off	Work On	everRun has powered off the PM because of an excessive failure rate.	
le.	X Booting	Finalize	PM is booting.	
ē	Rebooting	Finalize	PM is rebooting.	
8	Running	Finalize Shutdown Reboot Recover Replace	PM is running in Maintenance Mode. See Maintenance Mode.	

The everRun Availability Console

Using the everRun Availability Console

The Physical Machines Page

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. To learn more about the controls on this page, see Virtual Machine Actions or the help topic for a specific task.

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

- Viewing information about a VM, including its name, operating system, description, and resources in the tabs of the bottom pane
- Creating, copying, exporting, importing, or restoring VMs, as described in Creating and Migrating Virtual Machines
- Opening a Virtual Machine Console Session
- Reprovisioning Virtual Machine Resources
- Creating VM snapshots that can be restored or exported, as described in Creating a Snapshot
- Removing VM snapshots, as described in Removing a Snapshot
- Controlling the power state of a VM, as described in:
 - Starting a Virtual Machine
 - Shutting Down a Virtual Machine
 - Powering Off a Virtual Machine
- Removing a Virtual Machine or Renaming a Virtual Machine
- Performing advanced tasks or troubleshooting, as summarized in Advanced Topics (Virtual Machines)
- Mounting (and unmounting) a USB device or a network-mounted folder for use by the guest operating system, as described in Mounting a USB Device or Network-mounted Folder on the everRun System
- Attaching (and detaching) as USB device to a VM, as described in Attaching a USB Device to a Virtual Machine

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

Managing Virtual Machines

Using the everRun Availability Console

Virtual Machine Actions

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

Action	Description		
Za Create	Launches the Create VM Wizard. See Creating a New Virtual Machine.		
уу Сору	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.		
import/Restore	Imports a VM from a set of OVF and VHD files. See Creating and Migrating Virtual Machines. The import wizard allows you to <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. Open Virtual Machine Format (OVF) is an open standard for packaging and distributing physical or virtual machine data. The OVF format contains meta-data information about the VM. A Virtual Hard Disk (VHD) is a file that contains the virtual disk information.		
The following actions are available for use if the VM is running.			
	Mounts a USB device or a network-mounted folder (that is, a directory) to make it available to the guest operating system. You can then export a VM to the mounted location. See Mounting a USB Device or Network-mounted Folder on the everRun		
Mount	System.		

Action	Description	
Linmount	Unmounts a mounted USB device or a network-mounted folder. See Mounting a USB Device or Network-mounted Folder on the everRun System.	
Console	Opens a console for the selected VM. See Opening a Virtual Machine Console Session.	
I M Snapshot	Creates a VM snapshot that you can export to OVF and VHD files. See Managing Snapshots.	
Shutdown	Shuts down the selected VM. See Shutting Down a Virtual Machine.	
Power Off	Immediately stops processing in the selected VM and destroys its memory state. Use this only as a last resort, when the VM cannot be successfully shutdown. See Powering Off a Virtual Machine.	
The following action	s 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.	
یٹ Restore	Recovers an existing VM on your everRun system by overwriting the VM from a pre- vious backup copy of OVF and VHD files. See Replacing/Restoring a Virtual Machine from an OVF File.	
Export	Saves the image of a VM to a set of OVF and VHD files. You can import these files on another system or import them back to the same everRun system to restore or duplicate the original VM. See Exporting a Virtual Machine .	

Action	Description		
	Creates a VM snapshot that you can use to create a new VM or export to OVF and VHD files. See Managing Snapshots.		
Snapshot			
Þ	Boots the selected VM. See Starting a Virtual Machine.		
Start			
	Boots a VM from the selected virtual CD. See Booting from a Virtual CD.		
Boot From CD			
×	Removes a VM. See Removing a Virtual Machine.		
Remove			
The following action is available if the everRun software has removed the VM from service and powered it off because an excessive failure rate.			
•	Resets the mean time between failures (MTBF) counter for a VM so it can be brought back into service. See Resetting MTBF for a Failed Virtual Machine.		
	When a VM crashes, the everRun software automatically restarts it, unless it has		
Reset Device	fallen below its MTBF threshold. If the VM is below the MTBF threshold, the everRun		
	restart the VM and reset the MTBF counter.		

Managing the Operation of a Virtual Machine

The Virtual Machines Page

Using the everRun Availability Console

Virtual Machine States and Activities

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

State	Activity	Enabled Actions	Description
Q	 Installing 		The everRun software is installing the boot volume for a new VM.
	stopped	Start Copy Config Export Snapshot Boot From CD Remove	VM has been shut down or powered off.
C2	booting	Console Power Off	VM is starting.
*	running	Console Snapshot Shutdown Power Off	VM is operating normally on redundant physical machines
4	× running	Console Shutdown Power Off	VM is operating normally, but is not running on fully redundant resources.
Q	× stopping	Power Off Remove	VM is being shut down in response to the Shut- down action, or shut down because the remaining physical machine is transitioning into maintenance

State	Activity	Enabled Actions	Description
			mode.
Ø	* crashed		VM crashed and is restarting. If enabled, e-Alerts and support notification messages are sent.
\$	* crashed		VM crashed too many times and exceeded its MTBF threshold. The VM is left in a crashed state until Reset Device is clicked. See Resetting MTBF for a Failed Virtual Machine.

Managing the Operation of a Virtual Machine

The Virtual Machines Page

Using the everRun Availability Console

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.

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.

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
- Creating a Virtual Machine from a Snapshot

- Removing a Snapshot
- Adding a description for each volume, in the **Description** text box

Managing Snapshots

Using the everRun Availability Console

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.) 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).

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

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

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

Related Topics

The everRun Availability Console

Using the everRun Availability Console

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, size available, 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, sector size, 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 (^Q) 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.

Related Topics

The everRun Availability Console

Using the everRun Availability Console

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, link condition, name, internal name, type, number of connected Virtual Machines (VMs), speed, and MTU. You can also add a description for the network using the **Description** tab in the bottom pane.

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

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

- Connecting Additional Networks.
- Fixing a Network Connection.
- Viewing a list of the physical adapters that compose the network, on the **Summary** tab.
- Adding a description for a network, on the **Description** tab.
- Viewing a list of virtual machines that use the network, on the Virtual Machines tab.

- Changing the name by double-clicking the name in the **Name** column.
- Setting the MTU for A-Link and business networks.

For additional information on networks, see the following:

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

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

Related Topics

П

The everRun Availability Console

Using the everRun Availability Console

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 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 for details.

The Networks Page

The everRun Availability Console

Using the everRun Availability Console

Setting the MTU

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

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

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

Related Topics

The Networks Page

The everRun Availability Console

Using the everRun Availability Console

The Virtual CDs Page

Use the **Virtual CDs** page to create virtual CDs (VCDs). Use VCDs to make software installation or recovery media available to the virtual machines on the system. To open this page, click **Virtual CDs** in the left-hand navigation panel of the 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
- Removing a Virtual CD
- Renaming a Virtual CD
- Adding a description for each volume, in the **Description** text box

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

Related Topics

Inserting a Virtual CD

Ejecting a Virtual CD

Using the everRun Availability Console

The Upgrade Kits Page

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

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

For information about upgrading the everRun software, see Upgrading everRun Software.

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

Related Topics

The everRun Availability Console

Using the everRun Availability Console

Creating a USB Medium with System Software

You can use the **Upgrade Kits** page to create a USB medium with a copy of the install ISO file of the system software, everRun. You then use the USB medium to install the software on other nodes.

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

To create a USB medium with system software

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

The Create USB Medium dialog box opens.

If the node has more than one USB medium, you need to select a medium from the drop-down list.
 Then, click Continue (or click Cancel to cancel the procedure).

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

Use the USB medium to install the software on other nodes. See Software Installation.

Related Topics

The Upgrade Kits Page

The Users & Groups Page

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

To manage domain user accounts

For information about enabling the Active Directory service on your everRun system, see Configuring Active Directory. To grant or remove access for domain users to manage the everRun system, see Managing Domain User Accounts.

ñ

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 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

Managing Domain User Accounts

Managing Local User Accounts

Managing Local User Accounts

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

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
- VM Manager: ability to manage VMs (see Managing Virtual Machines for detailed information)
- Read-only: ability to view but not to change system configuration or to install system software

For the procedures below, begin by opening the **Users & Groups** page: in the left-hand navigation panel, select **Users & Groups**.

To add a user account

- 1. In the top pane, click Add.
- 2. In the Role drop-down window, select Administrator, Platform Manager, VM Manager, or Read-only.

- Provide values for the User Name, Password (and Confirm Password), Email Address, and Real Name fields. User names and passwords may be from 1 to 64 characters long, and must include no white space.
- 4. Click Save.

To edit a user account

- 1. Select the account you want to edit.
- 2. In the top pane, click Edit.
- 3. Change the user's information, as necessary. For example, to change a user's role, in the **Role** drop-down window, select **Administrator**, **Platform Manager**, **VM Manager**, or **Read-only**.
- 4. Click Save.

To remove a user account

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

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

Related Topics

П

Configuring Active Directory

Managing Domain User Accounts

The Users & Groups Page

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**.) 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.

For the procedures below, open the everRun - Grant Access Wizard:

- 1. In the left-hand navigation panel, click Users & Groups.
- 2. Click Grant Access.

To grant access to a domain user account

- 1. In the everRun Grant Access Wizard, specify the search range in the Search for menu.
- 2. Type the name or group for which to search. Partial names and text are allowed.
- 3. Click Search.
- 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.
- 5. 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.
 - VM Manager
 –Enables ability to manage VMs (see Managing Virtual Machines for detailed information)
 - Read Only–Enables read access but no management functions.
- 6. Click Finish. The new domain users are displayed in the Grant Access wizard.

To remove access for a domain user account

- 1. In the everRun Grant Access Wizard, click the check box next to users or groups you want to remove.
- 2. Click **Deny Access**, then click **Finish**.

Related Topic

Configuring Active Directory

4

Chapter 4: Upgrading everRun Software

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

- Use an upgrade kit. See Upgrading everRun Software Using an Upgrade Kit.
- Use a DVD. See Upgrading everRun From Release 7.3.4.0 with a DVD.

Related Topics

The Upgrade Kits Page

The everRun Availability Console

Using the everRun Availability Console

Upgrading everRun Software Using an Upgrade Kit

This topic describes how to use an upgrade kit of everRun software to upgrade the system software. The topic also explains how to download the kit and then upload it to the system, if you need to do so before upgrading the system. (If you want to upgrade a system with a DVD, see Upgrading everRun From Release 7.3.4.0 with a DVD.) You can optionally control the upgrade by enabling pauses. Inspecting a system during a pause is useful for verifying or reconfiguring third-party tools or other services that are not managed by the system.

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

Prerequisites:

- Before upgrading an everRun system, you should perform various system checks. For detailed information, access the Knowledge Base to search for the article *Pre-upgrade system check for everRun* (KB-4061). See Accessing Knowledge Base Articles.
- All PMs and VMs must be in good health before upgrading the system software. Before starting an upgrade, examine the everRun Availability Console to verify that there are no alerts indicating PM or VM problems.
- Eject any VCDs or USB media from the VMs before upgrading the system software. If VCD or USB media is still connected to the VMs, it prevents the system from migrating the VMs and putting the PMs into maintenance mode for the upgrade process.
- To verify that the system meets the requirements of the upgrade kit, use the **Qualify** button or the AVCLI kit-qualify command, as described in this topic.

Note: The upgrade also upgrades the AVCLI software on the 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 https://www.stratus.com/services-sup-port/downloads/?tab=everrun. For information about how to manually install AVCLI on a remote computer see AVCLI Command Overview.

The steps are:

٥

٥

- I. To download the upgrade kit
- II. To upload the upgrade kit to the system
- III. To qualify the software (optional)
- IV. To upgrade the system software

I. To download the upgrade kit

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

- 1. Open the Downloads page at https://www.stratus.com/services-support/downloads/?tab=everrun.
- 2. Scroll down to the upgrade section and then click the upgrade link to download the kit.

3. Navigate to a location on a local computer to save the file. If necessary, transfer the file to the remote management computer running the everRun Availability Console.

II. To upload the upgrade kit to the system

Upload the upgrade kit, if necessary, using one of the following methods:

- The Upgrade Kits page
 - a. In the everRun Availability Console, click Upgrade Kits in the left-hand navigation panel.
 - b. On the **Upgrade Kits** page, click the **Add a Kit** button beneath the masthead, which opens the **everRun Kit Upload Wizard**.
 - c. 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.
 - d. After you have selected a .kit file, click Upload, Import, or Finish (they perform the same function). A message such as Uploading file (DO NOT CLOSE WIZARD) appears while the file is uploading. The upload may require up to two minutes for a file stored locally, to ten or more minutes for a file stored over a network. If the upload fails, the wizard displays the message Failed to upload file.
 - e. After the upload is complete, the wizard closes and the Upgrade Kits page lists the state and version number of the upgrade kit. The Qualify, Upgrade, and Delete buttons also appear with the Add a Kit button.
 - f. If more than one upgrade kit is loaded, select the one to use.
- AVCLI commands— Add an upgrade kit by issuing the command avcli kit-add.

III. To qualify the software

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

Do so using one of the following methods:

- On the Upgrade Kits page, select the upgrade kit you want to qualify and then click Qualify.
- Issue the command avcli kit-qualify.

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

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

IV. To upgrade the system software

- 1. Begin the upgrade using one of the following methods:
 - On the Upgrade Kits page, click Upgrade.
 - Issue the command avcli kit-upgrade.

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

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

П

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

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

Note: After upgrading to everRun Release 7.5.0.5, shut down and then restart all running VMs, to enable features and performance improvements available to VMs in Release 7.5.0.5. Doing so immediately after the upgrade is not necessary, but shutting down and restarting VMs is a requirement for enabling VMs to run with the full capability available in Release 7.5.0.5. For information on shutting down and then starting VMs, see Managing the Operation of a Virtual Machine.

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

The Upgrade Kits Page

AVCLI Command Descriptions (see Kit Information)

The everRun Availability Console

Using the everRun Availability Console

Upgrading everRun From Release 7.3.4.0 with a DVD

You can upgrade everRun software from everRun Release 7.3.4.0 to a later release a using a DVD.

Prerequisites:

 \mathcal{O}

ñ

- Before upgrading an everRun system, you should perform various system checks. For detailed information, access the Knowledge Base to search for the article *Pre-upgrade system check for everRun* (KB-4061). See Accessing Knowledge Base Articles.
- All PMs and VMs in must be in good health before performing an upgrade of the software.
 Before starting an upgrade, examine the everRun Availability Console to verify that there are no alerts indicating PM or VM problems.
- Eject any VCDs or USB media from the VMs before upgrading the software. If VCD or USB media is still connected to the VMs, it prevents the system from migrating the VMs and putting the PMs into maintenance mode for the upgrade process.
- Upgrading from Release 7.3.4 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 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, contact your authorized Stratus service representative to access the Knowledge Base 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

- The everRun Availability Console
- Using the everRun Availability Console

Chapter 5: Managing Logical Disks

Manage logical disks using the everRun Availability Console. For overview information see Logical Disk Management as well as Logical Disks and Physical Disks.

To perform tasks, see the following:

- Responding to a Failed Logical Disk
- Activating a New Logical Disk
- Creating a New Storage Group
- Deleting a Storage Group
- Assigning a Logical Disk to a Storage Group

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.

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.

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

Related Topics

Logical Disks and Physical Disks

The everRun Availability Console

Storage Requirements

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.) 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. If you attempt to recover a missing or failed logical disk with the **Repair** button in the masthead of the everRun Availability Console, the system may be slow to repair the disk. Although the system successfully removes the failed logical disk from its storage group, it may be slow to migrate the data from the failed disk to other disks in the storage group. The **Alerts** page may continue to report that the logical disk is not present, that volumes have failed, and that storage is not fault tolerant. Also, the **Volumes** page may continue to show volumes in the broken (*) state. If this condition persists, contact your authorized Stratus service representative for assistance.
- 3. Repairing storage causes virtual machines (VMs) that are using failed logical disks to become simplex until repair is complete.
- 4. Systems configured for UEFI will only boot from the logical disk that the everRun software was originally installed on.
- 5. 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 ever-Run 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.

Related Topics

Logical Disks and Physical Disks

The everRun Availability Console

Activating a New Logical Disk

In an everRun system, the RAID controller creates logical disks from the system's physical disks. The ever-Run 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 noden 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.



Prerequisite: Note any relevant requirements in Storage Requirements.

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

Logical Disk Management

Logical Disks and Physical Disks

The everRun Availability Console

Storage Requirements

Creating a New Storage Group

You can create a new storage group to allocate additional storage space for virtual machines (VMs) and data. 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 for more information.

To create a new storage group

- 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

Deleting a Storage Group

Storage Requirements

Deleting a Storage Group

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

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

Assigning a Logical Disk to a Storage Group

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



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 Add to Storage Group (if the logical disk is unassigned) or Move to Storage Group (if the logical disk is currently assigned to another storage group).
- 5. In the dialog box that appears, click the **Storage Group** drop down box and select a storage group.
- 6. Click Add to Storage Group or 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
- Storage Requirements

Chapter 6: Managing Physical Machines

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

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

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

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

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

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

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

Maintenance Mode

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

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

- The 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.
- Avoid entering both PMs into maintenance mode at the same time. To keep VMs running, at least one PM must be up and running normally. (If you need to shut down the entire everRun system, see Shutting Down a Physical Machine.)

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

Related Topics

The everRun Availability Console Managing Physical Machines Physical Machines and Virtual Machines The Physical Machines Page The Virtual Machines Page

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.)

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.
- 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

The everRun Availability Console

Managing Physical Machines

The Physical Machines Page

Shutting Down a Physical Machine

Shut down a physical machine (PM), or node, to stop it from running when you need to service or replace the PM. Use the following procedure to shut down one and only one PM from the everRun Availability Console.

Cautions: 1. Using the following procedures 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. 2. Do not use the -f (force) option with the halt, poweroff, or reboot command of the host operating system of a PM. Doing so causes FT guests that are active on the same PM to hang. Instead, use the everRun Availability Console and maintenance mode to shut-down 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

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

- 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. 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 shut down, its activity is **Xoff (in Maintenance)**. You must manually restart the PM.

Related Topics

Maintenance Mode The everRun Availability Console Managing Physical Machines The Physical Machines Page

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. In response to the alert, click Load Balance in the masthead to initiate automatic load balancing of a VM.

The icon on the **Virtual Machines** page under **Current PM** column indicates VMs that will migrate imminently.

• manually place on node *N*. Advanced users can manually assign a preferred PM (node) for each individual VM, rather than relying on the automatic policy, if preferred.

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

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



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

Troubleshooting Physical Machines

The following topic describes troubleshooting procedures for PMs:

• Recovering a Failed Physical Machine

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

Recovering a Failed Physical Machine

Recover a physical machine (PM), or node, 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. Recovering a failed PM, though, is different from installing the software for the first time. The recovery preserves all data, but it recreates the /boot and root file systems, re-installs the everRun system software, and attempts to connect to the existing system. (If you need to replace the physical PM hardware instead of recovering the system software, see Replacing Physical Machines, Motherboards, NICs, or RAID Controllers.)

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

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

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

Prerequisites:

.

- 1. Determine which PM you need to recover.
- 2. Check that a monitor and keyboard are connected to the PM.
- 3. Check that Ethernet cables are connected from the PM your are replacing 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.
- 4. If you want to use DVD or USB media to install the system software on the replacement PM, obtain installation software for the release that the PM has been running by using one of the following methods:
 - Create a bootable USB medium on the Upgrade Kits page, as described in Creating a USB Medium with System Software.
 - Download an install ISO from your authorized Stratus service representative.
 - Extract an install ISO into the current working directory from the most recently installed 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

If you download or extract an install ISO, save it or burn it to a DVD or USB medium. See Obtaining everRun Software.

To recover a PM (with PXE boot installation)

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

- 1. In the everRun Availability Console, click Physical Machines in the left-hand navigation panel.
- Select the appropriate PM (node0 or node1) and then click Work On, which changes the PM's Overall State to Maintenance Mode and the Activity state to running (in Maintenance).

- 3. After the PM displays running (in Maintenance), click Recover.
- 4. When prompted to select the type of repair, click PXE PM Recover Preserve Data.



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

- 5. Click **Continue** to begin the recovery process. The system reboots the target PM in preparation for the system software reinstallation.
- As the PM reboots, enter the firmware (BIOS or UEFI) setup utility, and enable PXE boot (boot from network) for the priv0 NIC.

The recovery process continues with no user interaction, as follows:

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

You can monitor the progress of the software installation at the physical console of the target PM.

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

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

To recover a PM (with DVD/USB installation)

ñ

Use the following procedure to recover a PM by reinstalling the system software from a DVD or USB

medium.

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



- Click Continue to begin the recovery process. The system shuts down the target PM in preparation for the system software reinstallation.
- Insert the bootable media or mount the ISO image on the target PM, and then manually power on the PM.
- As the target PM powers on, enter the firmware (BIOS or UEFI) setup utility and set the Optical Drive or USB media as the first boot device.
- 8. Monitor the installation process at the physical console of the target PM.
- 9. At the **Welcome** screen, use the arrow keys to select the country keyboard map for the installation.
- 10. At the **Install or Recovery** screen, select **Recover PM**, **Join system: Preserving data** and press **Enter**.



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

11. 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:

- 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.
- 12. 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.

i

П

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.

- 13. 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.
- 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 redisplays until you enter valid information.

- 15. At this point, the software installation continues without additional prompts.
- 16. When the software installation is complete, the target PM reboots from the newly installed system software.
- 17. As the target PM boots, you can view its activity on the **Physical Machines** page of the ever-Run Availability Console. The **Activity** column displays the PM as **(in Maintenance)** after the recovery is complete.
- 18. If applicable, manually reinstall applications and any other host-level software, and reconfigure the PM to match your original settings.
- When you are ready to bring the target PM online, click Finalize to exit maintenance mode.
 Verify that both PMs return to the running state and that the PMs finish synchronizing.

Related Topics

Maintenance Mode Managing Physical Machines The everRun Availability Console The Physical Machines Page

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. To perform specific management tasks, see the following topics.

To manage the operational state of a VM, see:

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

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

To create or configure a VM, see:

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

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

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

To complete advanced tasks, see:

- Assigning a Specific MAC Address to a Virtual Machine
- Selecting a Preferred PM for a Virtual Machine
- Changing the Protection Level for a Virtual Machine (HA or FT)
- Configuring the Boot Sequence for Virtual Machines
- Resetting MTBF for a Failed Virtual Machine
- Locating a Dump File in a Virtual Machine

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

- Perform tasks of the available function buttons and actions on The Virtual Machines Page.
- View all available tabs on The Virtual Machines Page, except the Support tab.
- Perform tasks of the available function buttons and actions on The Snapshots Page.
- Create and delete VCDs from the The Virtual CDs Page.
- Expand volume container size, as described in Expanding a Volume Container on the everRun System; however, the VM Manager cannot expand a volume, as described in Expanding a Volume on the everRun System.

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

Planning Virtual Machine Resources

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

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

- Planning Virtual Machine vCPUs
- Planning Virtual Machine Memory

- Planning Virtual Machine Storage
- Planning Virtual Machine Networks

Planning Virtual Machine vCPUs

Allocate virtual CPUs (vCPUs) to assign computing resources to a virtual machine (VM) on the 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 a system is equal to the minimum of the number of hardware threads presented by either physical machine (PM) in the system. For example, in a system where one PM that has 4 cores and 2 threads per core (8 vCPUs) and a second PM (in that system) that has 8 cores and 2 threads per core (16 vCPUs), the total number of vCPUs available is 8 vCPUs (fewest threads of either PM).
- The number of vCPUs available for all VMs is equal to the total number of vCPUs available on the ever-Run system minus the number of vCPUs allocated to the everRun system software. (You set system vCPUs to 2 or 4, as described in Configuring System Resources.)
- 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.
- Windows-based VMs: If you change the number of assigned vCPUs from 1 to *n* or *n* to 1, after restarting the VM at the end of the reprovisioning process (see Reprovisioning Virtual Machine Resources), you must shut down and restart the VM a second time. This allows the VM to correctly reconfigure itself for Symmetric Multiprocessing (SMP). The VM displays odd behavior and is not usable until it is restarted.
- The System page of the everRun Availability Console (see The System Page) 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.
- By design, a VM displays its vCPU as an Intel Xeon Sandy Bridge E312xx with the base CPU speed
 of the host CPU, regardless of the system's actual CPU and actual CPU speed. For example, in a VM
 that is running a Windows operating system, the system properties utility displays the CPU as Sandy
 Bridge and the CPU speed as the base CPU speed even when the system's CPU is not Sandy Bridge
 and you have used a tool to increase or boost the CPU speed. For additional information, access the

Knowledge Base to search for the article *VM's vCPU reports as a Sandy Bridge with the base CPU clock speed* (KB-9913). See Accessing Knowledge Base Articles.

- 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

Π

System Requirements Overview

Creating and Migrating Virtual Machines

Managing Virtual Machine Resources

Planning Virtual Machine Memory

Allocate memory to assign physical memory to a virtual machine (VM) on your 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) 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).
 For example, if the total amount of memory is 32 GB, and you allocate 2 GB to the system software, there are 30 GB of memory available to the VMs.
- You can provision a single VM with memory up to the total amount of memory available to the VMs. Each VM consumes its requested amount of memory plus an additional 20% memory for overhead.
- The minimum memory allocation is 256 MB, but 64-bit operating systems require 600 MB or more. Be sure to verify the memory requirements for your guest operating systems.
- The System page of the everRun Availability Console (see The System Page) 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

Creating and Migrating Virtual Machines

Managing Virtual Machine Resources

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.

To conserve storage space in a volume container, you can remove older or obsolete snapshots as described in Removing a Snapshot. If necessary, you can also expand a volume container as described in Expanding a Volume Container on the everRun System.

Create separate boot and data volumes for each VM

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

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

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

Observe the maximum volume size

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

Related Topic

Storage Requirements Creating and Migrating Virtual Machines Managing Virtual Machine Resources

Planning Virtual Machine Networks

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

Creating and Migrating Virtual Machines

Managing Virtual Machine Resources

Creating and Migrating Virtual Machines

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

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

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

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

• Migrating a Physical Machine or Virtual Machine to a System

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

• Exporting a Virtual Machine

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

Managing Snapshots

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 system or export the snapshot to OVF and VHD files on a network share.

• Importing an OVF File

Use the everRun Availability Console to import OVF and VHD files from another everRun system to the everRun system.

• Replacing/Restoring a Virtual Machine from an OVF File

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

To migrate or import a system from an Avance or everRun MX system, see Migrating From Avance or ever-Run MX Systems for considerations, and then see one of the following topics to migrate or import the VMs, depending on your needs:

 Migrating a Physical Machine or Virtual Machine to a System (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 ever-Run system.

• Importing an OVF File from an everRun MX System (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 ever-Run system.

• Importing an OVF File from an Avance System (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 system.

Related Topics

Managing Virtual Machines

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.)

П

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. You must perform a different procedure to create a Windows Server 2003 VM.



To create a new VM

- On the Virtual Machines page (see The Virtual Machines Page), click Create to open the VM Creation Wizard.
- 2. On the Name, Description, Protection and OS page:
 - Type the Name and an optional Description for the VM as they will appear in the everRun Availability Console.

The VM name must meet the following requirements:

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

- 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 \\machine_URL\ShareName (for example, \\192.168.1.34\MyISO_Folder).
 - Remote ISO via NFS—The boot source is a remote ISO file, accessed through NFS. For Repository, enter the URL of the remote system in the format *nnn.nnn.nnn* (do not include http:// or https://).

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

- d. Click Next.
- 3. On the vCPUs and Memory page:
 - a. Specify the number of **vCPUs** and the amount of **Memory** to assign to the VM. For more information, see Planning Virtual Machine vCPUs and Planning Virtual Machine Memory.
 - b. Click Next.

- 4. On the **Volumes** page:
 - a. Type the Name of the boot volume as it will appear in the everRun Availability Console.
 - b. Type the 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 and Planning Virtual Machine Storage.
 - 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**.

Select a storage group that best supports the sector size of the volume that you want to create (see Planning Virtual Machine Storage). 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.)
- f. Click Next.
- On the Networks page, select the shared networks to attach to the VM (for more information, see Planning Virtual Machine Networks). You can also enable (or disable) the network and specify the MAC address. To continue, click Next.
- 6. On the Creation Summary page:
 - a. Review the creation summary. If you need to make changes, click **Back**.
 - b. If you want to prevent a console session from automatically starting to observe the software installation, deselect **Launch Console**.
 - c. To accept the VM as provisioned and begin the software installation, click Finish.

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

- 7. 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.
- 8. After you install the operating system, configure the additional resources and software necessary for production use, as described in:
 - Configuring Windows-based Virtual Machines
 - Configuring Linux-based Virtual Machines

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

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

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

Related Topics

Copying a Virtual Machine Renaming a Virtual Machine Removing a Virtual Machine Creating and Migrating Virtual Machines Managing Virtual Machine Resources Managing the Operation of a Virtual Machine

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.

ñ

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

- Create a bootable virtual CD (VCD) of the Windows Server 2003 media as described in Creating a Virtual CD.
- 2. Perform steps 1 through 7 of the procedure as described in Creating a New Virtual Machine.
- 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.

Related Topics

П

Creating and Migrating Virtual Machines

Managing the Operation of a Virtual Machine

Copying a Virtual Machine

Copy a virtual machine (VM) if you want to clone an existing VM on your 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.)

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.



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 Virtual Machines page (see The Virtual Machines Page), select the VM that you want to copy and click Shutdown.
- 2. When the VM has stopped, click Copy to open the copy wizard.
- 3. On the Name, Description, and Protection page:
 - 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 and Modes of Operation.

- c. Click Next.
- 4. On the vCPUs and Memory page:
 - a. Specify the number of **vCPUs** and the amount of **Memory** to assign to the VM. For more information, see Planning Virtual Machine vCPUs and Planning Virtual Machine Memory.
 - b. Click Next.
- 5. On the Volumes page, 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**.

Ensure that you select a **Storage Group** that supports the sector size of the volume you are copying (see Planning Virtual Machine Storage) and select the **Sector Size** that matches the source volume (the copy feature 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.

- 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. To continue, click Next.

- On the Networks page, activate the check box for each shared network that you want to attach to the VM.
- 7. 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 HA or 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

Managing Virtual Machine Resources

Managing the Operation of a Virtual Machine

Migrating a Physical Machine or Virtual Machine to a System

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

Procedures below describe how to migrate a PM or VM over the network: download the *P2V client* (virt-p2v) ISO file, boot the P2V client ISO file on the source PM or VM, and then use the client to configure, initiate, and monitor the secure network transfer from source side. No configuration steps are required on the system until

1

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

Notes:

Ö

- The migration process supports PMs or VMs running only the following operating systems:
 - CentOS/RHEL 6.x or 7.x.
 - Microsoft Windows 7, 8.x, or 10; or Windows Server 2008 R2, 2012, or 2016.
 Windows Server 2003–To migrate this VM, you must perform a different procedure.
 See Migrating a Windows Server 2003 VM to an everRun 7.x System.
 - Ubuntu 14.04 LTS, 16.04 LTS, or 18.04 Server-After migrating this VM, you need to perform additional procedures. See To complete the migration of an Ubuntu VM.
 - VMware Release 6.x
- For Windows-based VMs that support hibernation or fast startup mode, you must disable these features before the migration process. To fully disable hibernation or fast startup mode, see the instructions to recover from a migration that fails with the error message Failed to mount '/dev/sda1: Operation not permitted below in Troubleshooting.
- For Linux-based PMs or VMs, consider editing the /etc/fstab file before the migration process to comment out entries for data volumes and allow only the boot volume to mount. Because Linux-based VMs use different device names on the everRun system, a new VM may boot into single-user mode if it cannot mount volumes with their original device names. You can restore the /etc/fstab entries with the correct device names after the migration, as described below in Troubleshooting.
- When migrating a VMware VM, you must shutdown the VM using operating system shutdown commands in addition to powering it off from the VMware console. If you shutdown the VM using only the VMware console, the migration will fail.
- The source PM or VM must be offline for the duration of the migration process. Consider scheduling a planned maintenance period for the migration.
- The time required for the migration depends on the size and number of volumes on the source system as well as the network bandwidth between the source and the target system. For example, transferring a source system with one 20 GB boot volume over a 1Gb network may take about 30 minutes.

0

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

Prerequisite: 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.

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

To prepare for migrating a PM to the everRun system

- Download the P2V client ISO file from the Downloads page at https://www.stratus.com/services-support/downloads/?tab=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.
- If you want to verify the integrity of the ISO image, use the associated fciv checksum file and the Microsoft File Checksum Integrity Verifier (FCIV) executable file, if you have already downloaded the Microsoft file to your system. Otherwise, use the MD5 checksum hash function.

Using fciv

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

you downloaded.

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

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

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

Using MD5

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

CertUtil -hashfile path_to_file MD5

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

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

To prepare for migrating a VM to the everRun system

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

Using fciv

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

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

fciv -v -xml virt-p2v-n.n.n-n.nnnnnnn.n.el6.centos.xml
If the command succeeds (that is, it returns the message All files verified successfully), go to the next step. If the command fails, repeat the download.

Using MD5

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

CertUtil -hashfile path_to_file MD5

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

- 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.
- 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.)
- Click Test connection. If the P2V client connects to the everRun system, click Next to continue. A page appears with sections for Target properties, Fixed hard disks, and other settings.

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

 In the Target properties section, 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 ever-Run system.)

- 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 operation options, see Creating a New Virtual Machine and Modes of Operation.)
 - b. Next to **Output format**, select the disk image format, **raw** or **qcow2**. (The qcow2 format supports snapshots.)
- If you want to save debugging messages from the migration process, optionally select the Enable server-side debugging check box. (The debugging messages are included if you generate a diagnostic file for your authorized Stratus service representative, as described in Creating a Diagnostic File.)
- 9. Select which **Fixed hard disks** (volumes) to include in the migration by activating the check box next to each device.

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

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) 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 **Configure Network** dialog box and select the shared network from a drop-down list.

П

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

Click **OK** when you have finished configuring the network interface.

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

Click Finish on the last wizard page to implement the changes.

- 3. Click Start to boot the new VM.
- 4. Click **Console** to open the console of the VM and log on to the guest operating system. (For information about using the console, see Opening a Virtual Machine Console Session.)
- 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
- Diskeeper
- If you migrated from a VM source, disable any services associated with other hypervisors. Examples include:
 - VMware Tools
 - Hyper-V Tools
 - Citrix Tools for Virtual Machines

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

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

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

To complete the migration of an Ubuntu VM

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

After migrating an Ubuntu 16.04 VM

- 1. From the everRun Availability Console, open a console window into the VM.
- 2. Log in to the VM and go to the terminal.
- 3. Enter the following command: cd /etc/network.
- 4. Enter the following command: sudo vi interfaces.
- 5. Change the enol interface to ens3f0.
- 6. Save the changes.

7. Enter the following command: sudo systemctl restart networking.service.

After migrating an Ubuntu 14.04 VM

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

- 1. From the everRun Availability Console, open a console window into the VM.
- 2. Log in to the VM and go to the terminal.
- 3. Enter the following command: ifconfig.

Note that the command output does not include eth0.

4. Enter the following command: if config -a.

Note that the command output includes eth0.

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

In the interfaces file, change em1 to eth0.

- 7. Enter the following command: cd run.
- 8. Enter the following command: ls -1.
- 9. Enter the following command: sudo touch ifup.eth0.
- 10. Enter the following command: sudo vi ifstate.eth0.

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

- 11. Enter the following command: sudo if down eth0.
- 12. Enter the following command: sudo if up eth0.

The following procedure requires you to reboot the VM:

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

Troubleshooting

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

To cancel the migration process

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

To clean up after a canceled or failed migration

Open the 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. The diagnostic file includes any server-side debugging messages from the migration process.

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

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

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

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

- 1. Log on to the operating system of the source PM or VM.
- 2. Open the Power Options control panel and click Choose what the power buttons do.
- Next to When I press the power button, select Shutdown (instead of Hibernate or Sleep, if present).
- Under Shutdown Settings, clear the check box next to Turn on fast startup (recommended), if present.

- 5. Click Save changes.
- 6. Open Administrator Power Shell and execute the following command:
 - > powercfg /h off
- 7. Shut down the operating system and restart the migration process.

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

A Linux-based VM may fail to exit the **booting** state in 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 the 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 the 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.

To manually install a new network driver

After migrating a PM or VM, the network driver might not be properly installed (for example, Device Manager might list the driver with a warning, ⁽¹⁾). In this situation, manually install the driver:

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

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

Related Topics

Migrating From Avance or everRun MX Systems

Creating and Migrating Virtual Machines

Configuring Windows-based Virtual Machines

Configuring Linux-based Virtual Machines

Managing Virtual Machine Resources

Managing the Operation of a Virtual Machine

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.

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 (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 (In-Place Migration)

Use this procedure to perform an in-place migration of an everRun MX system and its VMs to the ever-Run 7.*x* software.

• Planning to Migrate from an Avance Unit (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 (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

П

Software Installation

Post-Installation Tasks

Creating and Migrating Virtual Machines

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 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.

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.

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 a System) or the import process (as described in Importing an OVF File from an everRun MX System).

Network Preparation

Prepare the platform network and the networking environment to conform to the everRun 7.*x* requirements. See General Network Requirements and Configurations.

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 recommend 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-topoint 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 <u>Storage Requirements</u>.

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 Software

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.

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.

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:

L

0

- 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 and Compatible Guest Operating Systems.

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

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

Software Installation

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

Migrating a Physical Machine or Virtual Machine to a System

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.
- Download the everRun 7.x ISO file from the Downloads page at https://www.stratus.com/services-support/downloads/?tab=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:

 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 ever-Run Availability Center. This is expected.

- 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 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. (You can option-ally 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.
- 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.

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:

П

- 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:

 If applicable, prepare your VMs for migration as described in Migrating a Physical Machine or Virtual Machine to a System. (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.)

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.

 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.
- In the left-hand navigation panel of XenCenter, locate and expand the entry for the everRun MX system. Click the VM, and click Start.
- 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 a System.
- 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 a System 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 Product License.

- 14. If necessary, configure and manage your VM as described in Managing Virtual Machines. For guest-specific settings, see:
 - Configuring Windows-based Virtual Machines
 - Configuring Linux-based Virtual Machines
- 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.
- Follow the instructions in Installing Software on the Second PM 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.

 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.

Notes:

П

- 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.
- 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 ever-Run MX system and click Save.
 - If you want to specify a static IP address for each node, click each Noden 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.

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 sys-

tem 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 ever-Run 7.*x* system.

For all systems, see Creating and Migrating Virtual Machines 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

П

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.

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.

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

Network Preparation

Prepare the platform network and the networking environment to conform to the everRun system requirements. See General Network Requirements and Configurations.

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.

Installing everRun Software

After the nodes in the everRun system are configured, you can install and configure the everRun software, as described in Software Installation.

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.

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:

0

- 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 and Compatible Guest Operating Systems.

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

Describes some items to consider when migrating or converting from an Avance unit to an everRun system.

Software Installation

Summarizes the steps for installing the everRun software.

Migrating a Physical Machine or Virtual Machine to a System

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.
- Download the everRun ISO file from the Downloads page at https://www.stratus.com/services-support/downloads/?tab=everrun.
- 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.
- 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

Π

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:

- 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.

П

5. Follow the instructions in Installing Software on the First PM 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. (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.
- Log on to the everRun Availability Console on node0 as described in Logging On to the everRun Availability Console for the First Time.

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:

Ö

- 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 ever-Run 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:

 If applicable, prepare your VMs for migration as described in Migrating a Physical Machine or Virtual Machine to a System. (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.)

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.
- 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 a System.
- 4. After the migration completes, power off the VM and close the VM console window.
- In the everRun Availability Console connected to the everRun node, verify that the VM appears on the Virtual Machines page.
- Start the migrated VM and verify that it is functioning properly. Follow the instructions in Migrating a Physical Machine or Virtual Machine to a System 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.

П

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 Product License.

- If necessary, configure and manage your VM as described in Managing Virtual Machines. For guest-specific settings, see:
 - Configuring Windows-based Virtual Machines
 - Configuring Linux-based Virtual Machines
- 8. Follow steps 1-7 to migrate additional VMs.
- 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**.
- Follow the instructions in Installing Software on the Second PM 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.)

- 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.

Notes:

ñ

- 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.
- 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 Noden 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.

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

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

Using fciv

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

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

fciv -v -xml virt-p2v-n.n.n-n.nnnnnnn.n.el6.centos.xml

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

Using MD5

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

CertUtil -hashfile path_to_file MD5

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

- 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 https://www.stratus.com/services-support/downloads/?tab=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.
- 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 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

- 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.

- 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.
- 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

I

Migrating From Avance or everRun MX Systems Creating and Migrating Virtual Machines Configuring Windows-based Virtual Machines

Managing Virtual Machine Resources

Managing the Operation of a Virtual Machine

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 a System.)

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:

0

- 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.
- 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 1Gb 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 sys-

i

tem, 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

 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**.
- 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.
- 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 https://www.stratus.com/services-support/downloads/?tab=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 work ing. 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

- 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.

0

- 4. Select To: Open Virtualization Format (OVF) Package. Click Next.
- 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.

 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.

ň

П

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.
- 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

 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 eachVHD 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.



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.
- On the Physical Machines page (see 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.
- On the Virtual Machines page (see The Virtual Machines Page), 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.

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.
- Update the VirtIO drivers to the latest supported versions, as described in Updating the VirtIO Drivers (Windows-based VMs).
- 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.
П

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.

[REVIEWERS-can you suggest example error messages/solutions?]

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

Creating and Migrating Virtual Machines

Configuring Windows-based Virtual Machines

Configuring Linux-based Virtual Machines

Managing Virtual Machine Resources

Managing the Operation of a Virtual Machine

Importing an OVF File from an Avance System

Import an Open Virtualization Format (OVF) file from an Avance unit if you want to transfer a VM to the ever-Run 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 a System.)

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.

1

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:

ð

- 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. 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 1Gb 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.

8

• 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.
- 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.
- 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:
 - Download the VirtIO.exe driver installation utility from the Drivers and Tools section of the Downloads page at https://www.stratus.com/services-support/downloads/?tab=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.
- On the Physical Machines page (see 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.
- On the Virtual Machines page (see The Virtual Machines Page), 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**.
- 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.
- 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.

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

П

0

- 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).
- 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 (Windowsbased 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.

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 Creating and Migrating Virtual Machines Configuring Windows-based Virtual Machines Configuring Linux-based Virtual Machines Managing Virtual Machine Resources Managing the Operation of a Virtual Machine

Importing an OVF File

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

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

This topic explains how to import an OVF file from a local computer, a USB device, or a remote file system such as an NFS export or a Windows share (also known as a CIFS share, such as, for example, Samba).. If you want to restore an existing VM on the same system to overwrite the VM and recover it from a previous

backup copy, see Replacing/Restoring a Virtual Machine from an OVF File. If you need to import an OVF file from an Avance system or an everRun MX system, see Importing an OVF File from an Avance System or Importing an OVF File from an everRun MX System.

Notes:

0

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

Prerequisite:

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

To import an OVF file

 \mathcal{O}

- 1. Log on to the everRun Availability Console on the target system.
- If you are importing a VM from a USB device or network share (instead of the PC running the everRun Availability Console), mount the device or share on the everRun system as described in Mounting a USB Device or Network-mounted Folder on the everRun System.
- 3. On the Virtual Machines page (see The Virtual Machines Page), click Import/Restore to open the Import/Restore Virtual Machine wizard.
- 4. Select one of the following:
 - Browse for OVF

Note: Browsing for VMware OVFs is not supported.

Click **Browse** and locate the OVF on a local computer.

Browse for OVF from USB

Select a partition from the pull-down menu, click **List OVFs**, and select a file from the pull-down menu.

Remote OVF via Windows Share(CIFS/SMB)

The OVF is located in a folder on a CIFS share. Enter values for **Username** and **Password**. For **Repository**, enter a value in the format **\\machine_URL\ShareName** (for example, **192.168.1.34\MyOVFsForImport**). Then, click **List OVFs** and select an OVF file from the list.

Remote OVF via NFS

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

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

Name, CPU, and Memory

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

Storage

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

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

- 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.
- 9. Click **Import** to begin importing the VM.

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

Note: Imported volumes begin to appear on the **Volumes** page of the 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.

 If applicable, use the Reprovision Virtual Machine wizard to allocate additional resources to the VM, as described in Reprovisioning Virtual Machine Resources. Also, if you want to allocate additional space in each volume container for snapshots, see Expanding a Volume Container on the everRun System.

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.
- For Windows-based VMs only, download and update the VirtIO drivers to the latest supported versions, as described in Updating the VirtIO Drivers (Windows-based VMs). (The correct VirtIO drivers are already present in Linux-based VMs.)

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

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 system may continue to synchronize data between PMs to enable High Availability (HA) or Fault Tolerant (FT) operation.

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

Troubleshooting

٥

П

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, if present.

To recover missing data volumes in the target VM

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

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

To recover missing network devices in the VM on the everRun system

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

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

To manually install a new network driver

After importing a PM or VM, the network driver might not be properly installed (for example, Device Manager might list the driver with a warning, ⁽¹⁾). In this situation, manually install the driver:

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

- 6. Select Red Hat VirtIO Ethernet Adapter.
- 7. Click **Next** to install the network driver.

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

Related Topics

Migrating From Avance or everRun MX Systems Mounting a USB Device or Network-mounted Folder on the everRun System Creating and Migrating Virtual Machines Configuring Windows-based Virtual Machines Configuring Linux-based Virtual Machines Managing Virtual Machine Resources Managing the Operation of a Virtual Machine

Replacing/Restoring a Virtual Machine from an OVF File

Replace a virtual machine (VM) from a everRun-created Open Virtualization Format (OVF) file if you want to restore (that is, recover) 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.)

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 (see Exporting a Virtual Machine) from an everRun system or a VM snapshot (see Exporting a Snapshot) to OVF and Virtual Hard Disk (VHD) files on a supported network share or a USB device. Copy these files to your management PC, or mount the USB device or network share on the target everRun system as described in Mounting a USB Device or Network-mounted Folder on the everRun System, and then use the everRun Availability Console on the target everRun system to restore the OVF and VHD files from your management PC.

1

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:

Ö

- You can restore a VM from only an OVF created from an everRun system. You cannot restore a VM from an OVF created from a third-party system.
- You typically restore a VM to recover the VM from a previous backup. When restoring a VM, the system attempts to preserve the hardware ID and MAC addresses of all network interfaces.
- Restore a VM only if you are specifically trying to restore a particular instance of an ever-Run VM and that the restored VM will be the only copy of this VM running across all ever-Run servers in your network.
- The time required to restore a VM depends on the size and number of volumes in the source VM as well as network bandwidth. For example, transferring a VM with one 20 GB boot volume over a 1Gb network may take about 30 minutes.
- If you overwrite and restore an existing VM, the 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). 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).
- 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.

Prerequisites:

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

To restore a VM

- 1. Log on to the everRun Availability Console on the target everRun system.
- If you are restoring a VM from a USB device or network share (instead of the PC running the everRun Availability Console), mount the device or share on the everRun system as described in Mounting a USB Device or Network-mounted Folder on the everRun System.
- 3. On the Virtual Machines page (see The Virtual Machines Page), select the VM that you want to restore in the upper panel.
- 4. In the lower panel, click **Restore** or click **Import/Restore** near the top pane.
- 5. Select one of the following:
 - Browse for OVF

Click **Browse** and locate the OVF on a local computer.

Browse for OVF from USB

Select a partition from the pull-down menu, click **List OVFs**, and select a file from the pull-down menu.

Remote OVF via Windows Share(CIFS/SMB)

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

Remote OVF via NFS

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

- 6. Select **Restore**. (Scroll down the window, if necessary.) A warning message appears.
- 7. Click **Continue** to proceed.



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

- 8. If prompted, add VHD files.
- 9. Review the information and make any desired edits, if necessary:

Name, CPU, and Memory

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

Storage

Displays the name, size, destination, and sector size of each volume. In the **Create** column, select a box for a volume to allocate a storage container for the volume on the everRun system (the boot volume is required). In the **Restore Data** column, select a 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) 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.

П

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.

- Optionally, clear the check box for Auto start Virtual Machine after restore if you need to reprovision the VM before starting it for the first time.
- 11. Click **Restore** to begin restoring the VM. When the transfer is complete, click **Done** to close the wizard.

Note: Restored volumes begin to appear on the **Volumes** page of the 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.

12. If applicable, use the **Reprovision Virtual Machine** wizard to allocate additional resources to the VM, as described in Reprovisioning Virtual Machine Resources. Also, if you want to allocate additional space in each volume container for snapshots, see Expanding a Volume Container on the everRun System.

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

٥

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, if present.

Related Topics

Creating and Migrating Virtual Machines Managing Virtual Machine Resources Managing the Operation of a Virtual Machine

Exporting a Virtual Machine

Export a virtual machine (VM) from a system in order to save an image of the VM to a network-mounted folder (that is, directory) or to a USB device. Exporting a VM from an 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. An exported VM can function as a backup of the original VM. You can directly export a VM from the everRun system as described in this topic, or you can create and export a snapshot. For an overview of snapshots, see Managing Snapshots.

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

Notes:

Ö

- Because the source VM must be shut down to export is, 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).
- The time required for the export depends on the size and number of volumes in the source VM as well as network bandwidth. For example, transferring a VM with one 20 GB boot disk over a 1Gb network may take about 30 minutes.
- If you will continue to use the source VM after exporting it, remember to set a different MAC address and IP address for the VM when you import it on the target system.
- If the 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.
- The maximum size of a vfat file that you can export is 4 GB. If you try to export a vfat file larger than 4 GB, the export will fail.
- For Linux-based VMs, when exporting a VM to another system, you do not need to modify the /etc/fstab file. When exporting a VM from an Avance system to an everRun system, consider editing the 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.



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.
- Click Console to open the console of the VM and log on to the Windows guest operating system.
- Ensure that all volumes are labeled accurately, as summarized in Managing Windows Drive Labels.
- 5. Run the Windows System Preparation Tool (Sysprep) to prepare the guest operating system for redeployment.

To export a VM

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

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

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

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

Mount device via Windows Share (CIFS/SMB)

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

Mount device via NFS

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

Mount USB

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

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

If the mount succeeds, the repository appears under **Device URL** and the **Export VM** button becomes active; otherwise, an alert appears.

- Select the volumes to include under Boot Volume to Export and Data Volumes to Export. (The boot volume is required.)
- 8. Click Export VM to export the VM.

You can monitor the **Export Status** in the **Summary** tab for the VM that you are exporting. Progress is reported as the percentage (%) completed for the whole export and for each volume. When the process is complete, the status changes to **Export completed successfully**. To cancel the export, click **Cancel** next to the **Export progress** percentage. A dialog box opens, asking you to confirm the cancellation. Click **Yes** to cancel.

The 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.

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

Troubleshooting

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

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

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

Related Topics

Attaching a USB Device to a Virtual Machine

Creating and Migrating Virtual Machines

Managing Virtual Machine Resources

Managing the Operation of a Virtual Machine

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

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

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



To mount a USB device or network-mounted folder

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

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

Mount device via NFS

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

Mount USB

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

4. Click Mount.

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

To unmount a USB device or network-mounted folder

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

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

Related Topics

Exporting a Virtual Machine

Managing Virtual Machines

Managing Windows Drive Labels

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

Configuring Windows-based Virtual Machines

Configuring Windows-based Virtual Machines

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

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

If you plan to create VM snapshots (see Managing Snapshots), 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)

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);
 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).

Configure the guest operating system to generate a crash dump file if the operating system crashes.
 Follow the instructions in the Microsoft article, *How to generate a complete crash dump file or a kernel crash dump file by using an NMI on a Windows-based system* (Article ID: 927069). Follow the instructions in the More Information section.

Related Topics

Managing Virtual Machines

Updating the VirtIO Drivers (Windows-based VMs)

Update the Red Hat VirtIO drivers in your Windows-based virtual machines (VMs) to the latest supported versions, to ensure the proper operation of the VMs. For example, you should update the VirtIO drivers after upgrading the system software (Upgrading everRun Software) or after using the P2V client to migrate a VM or a physical machine (PM) to the everRun system (Migrating a Physical Machine or Virtual Machine to a System).

To update the VirtIO drivers, first download them and create a VCD. Then, update the drivers in the VMs.

Notes:

Ö

- For proper operation, ensure that you download the VirtIO drivers only from the everRun Support page, as described in the following procedure. The VirtIO ISO file on the support page contains versions of the VirtIO drivers that have been tested with the everRun software, and they are known to work. VirtIO drivers from other sources could have compatibility issues.
- When updating the VirtIO drivers, use only the Browse my computer for the driver software option and select the specific folder or .inf file that applies to the guest operating system. If you use the Search automatically for updated driver software option or select only the top level of the VirtIO VCD, Windows might automatically install an incorrect driver.
- In some cases, the guest operating system requests a restart after drivers are updated. If so, restart the guest operating system.

To download the VirtIO drivers and create a VCD

- 1. Download the VirtIO ISO file from the **Downloads** page at https://www.stratus.com/services-support/downloads/?tab=everrun.
 - 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.

 If you want to verify the integrity of the ISO image, use the associated fciv checksum file and the Microsoft File Checksum Integrity Verifier (FCIV) executable file, if you have already downloaded the Microsoft file to your system. Otherwise, use the MD5 checksum hash function.

Using fciv

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

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

fciv -v -xml virtio-win-n.n.nn.xml

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

Using MD5

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

CertUtil -hashfile *path_to_file* MD5

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

Open the everRun Availability Console and create a VCD of the VirtIO ISO file (see Creating a Virtual CD).

To update the VirtIO drivers in a Windows-based virtual machine

- 1. Open the everRun Availability Console and insert the VCD into the Windows-based VM (see Inserting a Virtual CD).
- 2. In the VM console window, open **Device Manager** in the guest operating system.

The method to open Device Manager varies depending on the release of the guest operating system. One method is to open the Control Panel and select **Device Manager**. Another method is to open a search window and type **Device Manager**.

3. Expand **Network adapters** and locate the **Red Hat VirtIO Ethernet Adapter**. There may be more than one adapter present depending on the number of network interfaces in your VM.

If the **Red Hat VirtIO Ethernet Adapter** is not present, the VirtIO driver is not installed. Expand **Other devices** and locate the unknown **Ethernet Controller** device. Update the driver for this device.

- a. Right-click the Red Hat VirtIO Ethernet Adapter (or Ethernet Controller) and select Update Driver Software. Click Browse my computer for the driver software, specify the location of the VirtIO Ethernet driver (netkvm) for your guest operating system, and finish updating the driver. (For example, to update the driver in a Windows Server 2012 R2 guest, select the NetKVM\2k12R2\amd64\netkvm.inf file on the VirtIO VCD.)
- b. Repeat the driver update for each additional **Red Hat VirtIO Ethernet Adapter** (or **Ethernet Controller**) device.
- 4. Expand Storage controllers and locate the Red Hat VirtIO SCSI controller. There may be more than one controller present depending on the number of volumes in your VM. If the Red Hat VirtIO SCSI controller is not present, the VirtIO driver is not installed. Locate the unknown SCSI controller device, and update the driver for this device:
 - a. Right-click the Red Hat VirtIO SCSI controller (or SCSI controller) and select Update Driver Software. Click Browse my computer for the driver software, specify the location of the VirtIO SCSI driver (viostor) for your guest operating system, and finish updating the driver. (For example, to update the driver in a Windows Server 2012 R2 guest, specify the viostor\2k12R2\amd64\viostor.inf file on the VirtIO VCD.)
 - b. Repeat the driver update for each additional Red Hat VirtIO SCSI (or SCSI controller) device.



Caution: Although the device name is the **Red Hat VirtIO SCSI** controller, you must select the storage driver file that is labeled **viostor**, and not **vioscsi** (if present). Installing the **vioscsi** driver may crash your VM.

 If you intend to use the QEMU guest agent, as described in Installing the QEMU Guest Agent for Application-Consistent Snapshots (Windows-based VMs), also update the VirtIO serial driver that is associated with the guest agent; otherwise, skip to the next step.

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. Update the driver for this device: .

- a. Right-click the VirtIO Serial Driver and select Update Driver Software.
- b. 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.)
- 6. If applicable, restart the guest operating system to load the updated drivers.

Related Topics

Configuring Windows-based Virtual Machines

Creating and Migrating Virtual Machines

Managing the Operation of a Virtual Machine

Creating and Initializing a Disk (Windows-based VMs)

Create and initialize a disk to prepare it for partitioning into volumes in a Windows-based virtual machine.

To create and initialize a disk in a Windows-based virtual machine

- 1. Use the 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.
- 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

٥

Opening a Virtual Machine Console Session

Configuring Windows-based Virtual Machines

Creating and Migrating Virtual Machines

Managing the Operation of a Virtual Machine

Installing Applications (Windows-based VMs)

Install an application in a Windows-based virtual machine by doing one of the following:

- Download the installation program to the guest operating system as an executable file or ISO file.
- Mount a network drive that contains the installation program.
- Create and insert a Virtual CD (VCD) that contains the installation program. See Managing Virtual CDs.

Related Topics

Opening a Virtual Machine Console Session

Configuring Windows-based Virtual Machines

Creating and Migrating Virtual Machines

Managing the Operation of a Virtual Machine

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.

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:

L

Ö

- 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.
- 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.

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 https://www.stratus.com/services-support/downloads/?tab=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

Creating and Migrating Virtual Machines

Managing the Operation of a Virtual Machine
Configuring Linux-based Virtual Machines

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

- Creating and Initializing a Disk (Linux-based VMs)
- Installing Applications (Linux-based VMs)

If you plan to create VM snapshots (see Managing Snapshots), consider installing the QEMU guest agent, as described in:

• Installing the QEMU Guest Agent for Application-Consistent Snapshots (Linux-based VMs)

In addition, ensure that you configure the following settings:

- Disable hibernation (enabled by default in some cases) to prevent the guest operating system from going into a power-saving state.
- Configure the power button action in the guest operating system to shut down the guest (and not to hibernate it) to allow the Shutdown VM button in the 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.
- 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 ever-Run 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 to resolve the issue.

For more information about these settings, see your Linux documentation.

Related Topics

Managing Virtual Machines

Creating and Initializing a Disk (Linux-based VMs)

Create and initialize a disk to make it available for storing data in a Linux-based virtual machine.

To create and initialize a disk in a Linux-based virtual machine

- 1. In the everRun Availability Console, create a new volume in a storage group, as described in Creating a Volume in a Virtual Machine.
- 2. In the Linux-based virtual machine, use the volume management tool or edit files as needed to initialize and mount the volume. See your Linux documentation for complete instructions.

The disk device names for a Linux-based virtual machine are /dev/vda through /dev/vdh, not the standard /dev/sda through /dev/sdh. The 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 Configuring Linux-based Virtual Machines Creating and Migrating Virtual Machines Managing the Operation of a Virtual Machine

Installing Applications (Linux-based VMs)

Install an application in a Linux-based virtual machine by doing one of the following:

- Download the installation package to the guest operating system as an executable file or ISO file.
- Mount a network drive that contains the installation package.
- Create and insert a Virtual CD (VCD) that contains the installation package. See Managing Virtual CDs.

Related Topics

Opening a Virtual Machine Console Session Configuring Linux-based Virtual Machines Creating and Migrating Virtual Machines Managing the Operation of a Virtual Machine

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.

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 transactionintensive applications. In this case, taking a snapshot without preparing your applications results in a *crashconsistent* 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 gemu-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:

L

П

- 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 use, schedule a maintenance period for the installation.

Related Topics

Configuring Linux-based Virtual Machines

Creating and Migrating Virtual Machines

Managing the Operation of a Virtual Machine

Managing the Operation of a Virtual Machine

Manage the operation of a virtual machine as described in:

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

For additional information configuration and troubleshooting information, see Advanced Topics (Virtual Machines).

Starting a Virtual Machine

Start a virtual machine (VM) to boot the VM's guest operating system. You can also configure a starting mode for a VM, for when the everRun system boots.

To start a virtual machine

- 1. On the Virtual Machines page, select a VM.
- 2. Click Start in the bottom panel.

To configure a starting mode for a virtual machine, for when the system boots

- 1. On the Virtual Machines page, select a VM.
- 2. Click the **Boot** tab in the bottom panel.
- 3. For Auto Start Mode, select one of the following:
 - Last–Return the VM to its state when the system was shutdown: if the VM was running, the VM is restarted when the system boots; if the VM was stopped, the VM is not started when the system boots.
 - **On**-Start the VM when the system boots.
 - Off- Do not start the VM when the system boots.
- 4. Click Save.

Related Topics

П

Shutting Down a Virtual Machine

Powering Off a Virtual Machine

Managing the Operation of a Virtual Machine

Shutting Down a Virtual Machine

Shut down a virtual machine (VM) to begin an orderly shutdown of the guest operating system.

Note: You can shut down a VM with guest operating system commands. Some guests allow (or can be configured to allow) you to shut down a VM using the everRun Availability Console.

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, the acpid package that enables the power button is not included in the default installation. You can manually install this package to enable the power button using the following command (or see the documentation for your guest operating system):

sudo apt-get install acpid

For versions of Ubuntu running the desktop, the everRun Availability Console **Shutdown** button causes the VM's Ubuntu desktop to prompt you with selecting one of three icons: suspend, sleep, or shutdown. To allow the Ubuntu VM to shutdown without the desktop prompts, you must modify the powerbtn file.

To modify the powerbtn file

- 1. In the VM, edit the /etc/acpi/events/powerbtn file.
- 2. Comment out these lines:

event=button[/]power
action=/etc/acpi/powerbtn.sh

3. Add these lines:

```
event=button/power (PWR.||PBTN)
action==/sbin/poweroff
```

4. Issue the following command to restart acpid:

```
systemctl restart acpid
```

See the documentation for your guest operating system to configure the behavior of the system power button, thus enabling the **Shutdown** button to work in the everRun Availability Console.

To shut down a VM in everRun Availability Console

- 1. On the Virtual Machines page, select a VM.
- 2. Click **Shutdown** in the bottom panel.

A warning message appears, asking you to confirm the shutdown. Click **Yes** to shutdown or **No** to discontinue the shutdown.

If the VM is not responding, you can also **Power Off** the VM to stop it without properly shutting down the guest operating system.

Related Topics

Starting a Virtual Machine

Powering Off a Virtual Machine

Managing the Operation of a Virtual Machine

Powering Off a Virtual Machine

Power off a virtual machine (VM) to stop it without properly shutting down guest operating system.

Caution: Use the **Power Off** command only if the **Shutdown** command or guest operating system commands fail. Powering off a VM is similar to pulling the power cord, which may result in data loss.

To power off a virtual machine

- 1. On the Virtual Machines page, select a VM.
- 2. Click **Power Off** in the bottom panel.

Related Topics

Starting a Virtual Machine

Shutting Down a Virtual Machine

Managing the Operation of a Virtual Machine

Advanced Topics (Virtual Machines)

Opening a Virtual Machine Console Session

Open a virtual machine (VM) console session to display the console of the guest operating system running in the VM.

The following procedure describes how to open a VM console session in the everRun Availability Console, but you can also use a remote desktop application for this purpose.

To open a VM console session

- 1. On the Virtual Machines page, select a VM.
- 2. Ensure that the VM is in a running state.
- 3. Click **Console** (**I**) in the bottom panel.

Troubleshooting

To resolve an issue where the VM console window does not open

Ask your network administrator to open ports 6900-6999 (inclusive).

To resolve an issue where the VM console window is blank

Verify that the VM is powered on and not in the process of booting. Also, click in the console window and press any key to deactivate the screen saver.

To resolve an issue where more than one VM console window is displayed and they are behaving erratically

Close all console windows and open only one console window.

To resolve an issue where the VM console window hangs on the 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 $\rm e$ 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.
- In the following line, replace vt100-nav with vt100 by deleting -nav at the end of the line. The updated line appears as follows:

Run gettys in standard runlevels co:2345:respawn:/sbin/agetty xvc0 9600 vt100

3. Save the inittab file.

Related Topics

Starting a Virtual Machine

Shutting Down a Virtual Machine

Managing the Operation of a Virtual Machine

Renaming a Virtual Machine

Rename a virtual machine (VM) to change its name as it appears on the Virtual Machines page.

If you need to change the host name of the guest operating system running in a VM, use guest operating system tools.



Prerequisite: To rename a VM, you must shut it down.

To rename a virtual machine

- 1. On the Virtual Machines page, select a VM.
- 2. Click **Shutdown** and wait for the VM to shut down.
- 3. Double-click the name of the VM.
- 4. Type the new name. The VM name must meet the following requirements:
 - A VM name must start with a word or a number, and the name cannot include the special characters (for example, #, %, or \$).
 - A VM name cannot use hyphenated prefixes such as Zombie- or migrating-.
 - A VM name has a maximum of 85 characters.
- 5. Press Enter.

Related Topics

Removing a Virtual Machine

Creating and Migrating Virtual Machines

Managing the Operation of a Virtual Machine

Removing a Virtual Machine

Remove a virtual machine (VM) to permanently delete it and optionally delete associated volumes from the everRun system.

Notes:

Ö

- 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.
- 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** in the bottom panel.
- 3. When the VM has stopped, click Remove.
- 4. In the **Remove Virtual Machine** dialog box, activate the check box next to volumes that you want to delete. Clear the check box for volumes to save as archives or save for attachment to another VM.

Caution: Make sure that you select the correct VM and volumes for removal. When you click **Delete VM**, these items are permanently removed.

5. Click **Delete VM** to permanently delete the VM and any selected volumes.

Related Topics

Renaming a Virtual Machine

Creating and Migrating Virtual Machines

Managing the Operation of a Virtual Machine

Managing Virtual Machine Resources

Manage virtual machine resources to reconfigure the vCPUs, memory, storage, or network resources of an existing virtual machine.

To reconfigure virtual machine resources, use the **Reprovision Virtual Machine** wizard, as described in:

• Reprovisioning Virtual Machine Resources

To reconfigure virtual machine volumes, see the following task-specific topics:

- Creating a Volume in a Virtual Machine
- Attaching a Volume to a Virtual Machine
- Detaching a Volume from a Virtual Machine
- Removing a Volume from a Virtual Machine
- Expanding a Volume Container on the everRun System
- Expanding a Volume on the everRun System

To recover virtual machine resources, freeing space for new volumes or virtual CDs, see:

• Recovering Virtual Machine Resources

To enable or disable virtual machine components, see:

• Enabling and Disabling VM Components

Reprovisioning Virtual Machine Resources

Reprovision a virtual machine (VM) to change its allocation of virtual CPUs (vCPUs), memory, storage, or network resources.

Launch the **Reprovision Virtual Machine** wizard by clicking **Config** in the bottom pane of the **Virtual Machines** page. The wizard steps you through the process of reallocating resources to the VM.

Prerequisites:

- Review the prerequisites and considerations for allocating vCPUs, memory, storage, and network resources to the VM, as listed in Planning Virtual Machine Resources. For more information about storage resources, see Planning Virtual Machine Storage.
- To reprovision a VM, you must shut down the VM.

To reprovision a virtual machine

- 1. Open the Virtual Machines page (see The Virtual Machines Page).
- 2. Select a VM and click Shutdown.

- 3. When the VM has stopped, click Config to display the Reprovision Virtual Machine wizard.
- 4. On the Name, Description, and Protection page:
 - Type the Name and an optional Description for the VM as they will appear in the everRun Availability Console

The VM name must meet the following requirements:

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

- c. Click Next.
- 5. On the vCPUs and Memory page:
 - Specify the number of vCPUs and the amount of Memory to assign to the VM. For more information, see Planning Virtual Machine vCPUs and Planning Virtual Machine Memory.
 - b. Click Next.
- 6. On the Volumes page, you can:

Notes:

Ö

- You cannot modify the VM boot volume, only data volumes. However, you can detach the boot volume.
- To expand a volume container, see Expanding a Volume Container on the everRun System.
- Click **Boot Volume** to detach the boot volume.

Caution: If you detach the boot volume, the VM becomes unbootable.

A warning appears saying that detaching the boot value causes the VM to become unbootable. If you want to undo detaching the boot volume, click **Undo Detach**.

- Click **Detach** to disconnect a volume from a VM and keep it for future use.
- Click **Delete** to permanently remove a volume from the everRun system.
- Select an unattached volume from a pulldown menu (if displayed) and click Attach.

You can also, if applicable, click **Add New Volume** to create a new data volume. (If the button is not visible, scroll down to the bottom of the wizard page.)

For an unattached volume or a new volume, specify the volume's parameters:

- a. Type the Name of the volume.
- b. Type the 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 and Planning Virtual Machine Storage.
- c. Select the Storage Group for the volume, and, if applicable, select the Sector Size.

Select a storage group that best supports the sector size of the volume (see Planning Virtual Machine Storage). 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 Attach to connect a volume to a VM.

To continue, click Next.

 On the Networks page, activate the check box for each shared network that you want to attach to the VM.

For each shared network that you attach, you can also optionally:

- Set a custom MAC address (for details, see Assigning a Specific MAC Address to a Virtual Machine).
- Set the State to Enabled or Disabled, which allows you to allow or block network traffic to the selected network.

For more information, see Planning Virtual Machine Networks. To continue, click Next.

8. On the Configuration Summary page:

Caution: Make sure that any volumes marked for removal are correct. When you click **Fin-ish**, permanent data loss occurs on disks marked for removal.

- a. Review the configuration summary. If you need to make changes, click **Back**.
- b. To accept the VM as provisioned, click Finish.
- 9. Click **Start** to restart the VM.
- 10. For Windows-based VMs, if you changed the number of assigned virtual CPUs in a Windows-based VM from 1 to *n* or *n* to 1, after restarting the VM at the end of the re-provisioning process, you must shut down and restart the VM a second time. This allows the VM to correctly reconfigure itself for Symmetric Multiprocessing (SMP). The VM displays odd behavior and is not usable until it is restarted.

Related Topics

Managing Virtual Machine Resources

Planning Virtual Machine Resources

Managing Virtual Machines

Creating a Volume in a Virtual Machine

Create a volume to attach a new, blank volume to a virtual machine (VM). (You can also attach an existing, unattached volume as described in Attaching a Volume to a Virtual Machine.)

Prerequisite: Before creating a volume for a VM, you must shut down the VM.

To create a new volume in a VM

- 1. Open the Virtual Machines page (see The Virtual Machines Page).
- 2. Select a VM and click **Shutdown**.
- 3. When the VM has stopped, click **Config** to display the **Reprovision Virtual Machine** wizard.
- 4. Click **Next** on each wizard page until the **Volumes** page is displayed. (If applicable, see Reprovisioning Virtual Machine Resources to configure additional VM resources.)
- 5. On the **Volumes** page, click **Add a new volume**. (If the button is not visible, scroll down to the bottom of the wizard page.)

- 6. Under **To Be Created**, do the following:
 - a. Type the Name of the volume as it will appear in the 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 and Planning Virtual Machine Storage.
 - c. Select the Disk Image format:
 - RAW raw disk format
 - QCOW2 QEMU Copy On Write (QCOW2) format, which supports snapshots

(For information on references to QCOW2, see Important Considerations.)

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). 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 guest operating system, as described in:
 - Creating and Initializing a Disk (Windows-based VMs)
 - Creating and Initializing a Disk (Linux-based VMs)

Related Topics

Detaching a Volume from a Virtual Machine Removing a Volume from a Virtual Machine Managing Virtual Machine Resources Planning Virtual Machine Resources Managing Virtual Machines

Attaching a Volume to a Virtual Machine

0

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

- Ensure that the volume you want to attach is not in use by another virtual machine; otherwise, you cannot attach it. Open the Volumes page, locate the volume, and ensure that the value in the Used By column is None.
- 2. Open the Virtual Machines page (see The Virtual Machines Page).
- 3. Select a VM and click **Shutdown**.
- 4. When the VM has stopped, click **Config** to display the **Reprovision Virtual Machine** wizard.
- 5. Click **Next** on each wizard page until the **Volumes** page is displayed. (If applicable, see Reprovisioning Virtual Machine Resources to configure additional VM resources.)
- 6. On the **Volumes** page, locate the pulldown menu next to the **Add a new volume** button. Select an unattached volume from the pulldown menu and click **Attach**.

(If the pulldown menu is not visible, scroll down to the bottom of the wizard page. The pulldown menu is displayed only if there are unattached volumes on the 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 Detaching a Volume from a Virtual Machine Removing a Volume from a Virtual Machine Managing Virtual Machine Resources Planning Virtual Machine Resources Managing Virtual Machines

Detaching a Volume from a Virtual Machine

Detach a volume to disconnect it from a virtual machine and keep it for future use, or attach it to another virtual machine as described in Attaching a Volume to a Virtual Machine. (You can also permanently delete the volume from the everRun system, as described in Removing a Volume from a Virtual Machine.)

Notes:

Ö

- 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.
- If you detach a boot volume from a VM, you cannot boot the VM; however, you might want to detach the boot volume to diagnose a boot problem or data corruption in the volume. You can temporarily attach the boot volume to another VM as a data volume, as described in Attaching a Volume to a Virtual Machine. After using guest operating system tools to resolve the issue, detach the volume and reattach it to its original VM.

Prerequisite: Before detaching a volume from a virtual machine, you must shut down the virtual machine.

To detach a volume from a virtual machine

- 1. Open the Virtual Machines page (see The Virtual Machines Page).
- 2. Select a VM and click **Shutdown**.
- 3. When the VM has stopped, click Config to display the Reprovision Virtual Machine wizard.
- 4. Click **Next** on each wizard page until the **Volumes** page is displayed. (If applicable, see Reprovisioning Virtual Machine Resources to configure additional VM resources.)
- 5. On the **Volumes** page, locate the volume to detach. (If the volume is not visible, scroll down on the wizard page.)
- 6. Click Detach beside the volume name to mark the volume for detachment.



- 7. Click **Next** on each wizard page until the **Configuration Summary** page is displayed. Verify the configuration changes.
- 8. Click **Finish** to detach the selected volume.

Related Topics

Attaching a Volume to a Virtual Machine

Removing a Volume from a Virtual Machine

Managing Virtual Machine Resources

Planning Virtual Machine Resources

Managing Virtual Machines

Removing a Volume from a Virtual Machine

Remove a virtual machine (VM) volume to permanently delete it from the 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.)

Notes:

0

- 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.
- 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).
- 2. Select a VM and click **Shutdown**.
- 3. When the VM has stopped, click Config to display the Reprovision Virtual Machine wizard.
- 4. Click **Next** on each wizard page until the **Volumes** page is displayed. (If applicable, see Reprovisioning Virtual Machine Resources to configure additional VM resources.)
- 5. On the **Volumes** page, locate the volume to delete. (If the volume is not visible, scroll down on the wizard page.)
- 6. Click **Delete** beside the volume name to mark the volume for deletion.

Caution: Be careful to mark the correct volume to remove, avoiding any volumes that are currently in use.

- 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

I



Caution: Before removing a volume, ensure that it is no longer needed by other

edministrators.

- 1. Open the **Volumes** page.
- 2. Select an unattached volume. (The **Used By** column must read **None**, otherwise, the **Remove** button is not displayed.)
- 3. Click **Remove**.

Related Topics

Detaching a Volume from a Virtual Machine

Attaching a Volume to a Virtual Machine

Managing Virtual Machine Resources

Planning Virtual Machine Resources

Managing Virtual Machines

Renaming a Volume on the 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 Detaching a Volume from a Virtual Machine Removing a Volume from a Virtual Machine Managing Virtual Machine Resources Planning Virtual Machine Resources Managing Virtual Machines

Expanding a Volume 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.)

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.

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

ñ

- 1. On the **Physical Machines** page (see 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.
- 2. On the Volumes page (see The Volumes Page), and select the volume that you want to expand.
- 3. In the bottom pane, click the **Container** tab and then click **Expand Container**.
- 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.

5. 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 Creating a Volume in a Virtual Machine Detaching a Volume from a Virtual Machine Removing a Volume from a Virtual Machine Managing Virtual Machine Resources Planning Virtual Machine Resources

Managing Virtual Machines

Expanding a Volume on the 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, 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

 If applicable, expand the volume container for the volume as described Expanding a Volume Container on the everRun System. 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.

- 2. On the **Virtual Machines** page (see The Virtual Machines Page), select the VM that contains the volume that you want to expand. Ensure that the VM is **stopped**.
- 3. In the bottom pane, click the **Volumes** tab and select the volume that you want to expand. In the **Action** column, click **Expand Volume**.
- 4. Next to Expand By, type the amount of storage space to add to the volume (in gigabytes (GB)). When you type the number, the dialog box displays the Expanded Volume Size that will result if you complete the operation.

Note: Consider the Expand By entry carefully, because after expanding a volume, you cannot undo the change or reduce the size of the volume; you can only expand the volume further.

5. Click **Expand Volume** to commit the change and expand the volume. The dialog box displays the expansion progress and automatically closes when the operation is complete.

Related Topics

Expanding a Volume Container on the everRun System

Creating a Volume in a Virtual Machine Detaching a Volume from a Virtual Machine Removing a Volume from a Virtual Machine Managing Virtual Machine Resources Planning Virtual Machine Resources Managing Virtual Machines

Recovering Virtual Machine Resources

To conserve storage space, remove VM resources when they are no longer needed. You may also need to immediately recover storage space when there is insufficient space for certain tasks, such as creating a volume or VCD.

To recover storage space, remove unused resources as described in the following topics:

- Removing a Virtual Machine
- Removing a Volume from a Virtual Machine
- Removing a Virtual CD

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

Related Topics

Managing Virtual Machine Resources Planning Virtual Machine Resources

Managing Virtual Machines

Enabling and Disabling VM Components

You can enable or disable certain virtual machine (VM) components on individual nodes using the everRun Availability Console. You can enable or disable volumes and networks on node1 and/or node0. You can also enable or disable instances of a VM on node1 or on node0.

Notes: You cannot enable (or disable) the same component on both nodes at the same time. You cannot disable both instances of a VM. You must follow an order of last-in, first-out (LIFO) when enabling VM volumes on both nodes. For example, if you disable a volume on node0 and then disable the volume on node1, and you then want to enable the volume on node0, you must first enable the volume on node1 before you can enable the volume on node0.

To enable or disable a virtual machine component

- 1. Open the Virtual Machines page (see The Virtual Machines Page).
- 2. Select a VM and then click the Support tab in the bottom panel.
- 3. Beneath the **Support** tab, locate the component that you want to enable or disable: **Volume**, **Network**, or **VM instance**.
- 4. In the row for the specific volume, network, or VM instance that you want to enable or disable, select **Enable node0** or **Enable node1**, or select **Disable node0** or **Disable node1**.
- 5. A **Confirm** dialog box appears, asking you to confirm the change. Click **Yes** to enable or disable the component.

Related Topics

Managing Virtual Machine Resources

Managing Virtual Machines

Managing Virtual CDs

Create and manage virtual CDs (VCDs) to make software installation media available to the virtual machines on your 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.

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

L

- Inserting a Virtual CD
- Ejecting a Virtual CD
- Booting from a Virtual CD
- Renaming a Virtual CD
- Removing a Virtual CD

Users who are assigned the role Administrator or Platform Manager can perform all VCD tasks. Users who are assigned the role VM Manager can perform all VCD tasks, except rename a VCD. (For information on assigning these roles, see Managing Local User Accounts.)

Creating a Virtual CD

Create a virtual CD (VCD) to make software installation media available to the virtual machines (VM) on your 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) to install a guest operating system or start a VM from a bootable recovery VCD. You can also insert a VCD into a running VM (see Inserting a Virtual CD) to install software applications.

Notes:

- 1. 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.
- 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

0

- 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 displays progress of the upload.

You can determine the status of a VCD by checking the State column on the Virtual CDs page:

- A syncing icon (⁽¹⁾) indicates that the VCD is still being created.
- A broken icon (X) indicates that the VCD creation failed. Remove the VCD and try creating it again.
- A normal icon () indicates that the transfer is complete and that the VCD is ready to use.

Related Topics

Inserting a Virtual CD Ejecting a Virtual CD Managing Virtual CDs Creating and Migrating Virtual Machines

Inserting a Virtual CD

L

٥

Insert a virtual CD (VCD) in a virtual machine (VM) to access installation media when installing applications in a guest operating system. (To attach a USB device, see Attaching a USB Device to a Virtual Machine. To boot a virtual machine from a VCD, see Booting from a Virtual CD.)

Caution: When you insert a VCD into a running, 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.

Note: By default, VCDs are enabled for insertion in VMs. To change this configuration, see Configuring VM Devices.

To connect a VCD to a VM

- 1. If necessary, create a VCD (see Creating a Virtual CD) for the software installation media you need to access.
- 2. On the Virtual Machines page, select a VM.
- 3. In the bottom pane, click the CD Drives & USB Devices tab.
- 4. To select a VCD, click Insert a CD and select a VCD. Use the pulldown menu, if it exists.

When the system has inserted the VCD, its name appears to the right of CD-ROM.

Related Topics

Creating a Virtual CD

Ejecting a Virtual CD

Booting from a Virtual CD

Managing Virtual CDs

Ejecting a Virtual CD

Eject a virtual CD (VCD) to disconnect it from a virtual machine (VM). Ejecting a VCD allows you to insert another VCD into the VM. It also makes the VCD available for inserting into another VM.

To eject a VCD from a VM

- 1. Unmount the VCD from the guest operating system to ensure that it is not in use.
- 2. On the Virtual Machines page, select a VM.
- 3. Click the CD Drives & USB Devices tab in the lower frame.
- 4. On the CD Drives tab, click Eject CD.

Related Topics

Creating a Virtual CD

Inserting a Virtual CD

Booting from a Virtual CD

Managing Virtual CDs

٥

Booting from a Virtual CD

Boot a virtual machine from a virtual CD (VCD) to install a guest operating system or to perform maintenance.

Before booting from a VCD, you must shut down the virtual machine.

To boot a virtual machine from a VCD

- 1. If necessary, create a VCD from a bootable CD/DVD (see Creating a Virtual CD).
- 2. On the Virtual Machines page, select a virtual machine.
- 3. If the virtual machine is running, click **Shutdown**.
- 4. When the virtual machine status shows stopped, click Boot from CD in the lower pane.
- 5. Select the bootable VCD, then click **Boot**.

Note: A Windows-based virtual machine booted from a VCD boots as a hardware virtual machine (HVM), and it can access only the first three disk volumes.

Related Topics

Creating a Virtual CD

Inserting a Virtual CD

Ejecting a Virtual CD

Managing Virtual CDs

Creating and Migrating Virtual Machines

Managing the Operation of a Virtual Machine

Renaming a Virtual CD

Rename a virtual CD (VCD) to change its name as it appears on the Virtual CDs page.

To rename a VCD

- 1. Locate the VCD on the Virtual CDs page.
- 2. Double-click the name of the VCD.
- 3. Specify the new name and press Enter.

Related Topics

Removing a Virtual CD

Inserting a Virtual CD

Ejecting a Virtual CD

Creating a Virtual CD

Managing Virtual CDs

Removing a Virtual CD

Remove a virtual CD (VCD) to permanently delete it from the 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

Inserting a Virtual CD

Ejecting a Virtual CD

Creating a Virtual CD

Managing Virtual CDs

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:

L

Ö

- 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
- Creating a Virtual Machine from a Snapshot
- Exporting a Snapshot
- Removing a Snapshot

Users who are assigned the role Administrator, Platform Manager, or VM Manager can perform these tasks. (For information on assigning these roles, see Managing Local User Accounts.)

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.

To view the snapshots that you have created in the everRun Availability Console:

- Open the **Snapshots** page (see The Snapshots Page)
- On the Virtual Machines page (see The Virtual Machines Page), 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. 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)
- Installing the QEMU Guest Agent for Application-Consistent Snapshots (Linux-based VMs)

Related Topics

Creating and Migrating Virtual Machines

Managing the Operation of a Virtual Machine

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. For an overview of snapshots, see Managing Snapshots.

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)
- Installing the QEMU Guest Agent for Application-Consistent Snapshots (Linux-based VMs)

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. If necessary, you can expand a volume container as described in Expanding a Volume Container on the 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:

ð

- 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)

- 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).
- Ensure that all volumes are labeled accurately, as summarized in Managing Windows Drive Labels.
- 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).

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), 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.
- 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. If you want to export a completed snapshot, see Exporting a Snapshot.

Related Topics

Managing Snapshots Creating and Migrating Virtual Machines Managing the Operation of a Virtual Machine

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. For an overview of snapshots, see Managing Snapshots.)

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:

Ö

- 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.
- 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.
- 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), 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.
- 4. On the Snapshots page, select the snapshot to use as a source for the new VM.

Snapshots are usually in a normal state (\checkmark) the **State** column. If a snapshot is broken (\thickapprox), it may indicate that one or more volumes in the snapshot are unavailable for the VM creation.

- 5. In the bottom pane, click Create VM.
- 6. The **Create VM from Snapshot "name"** dialog box appears with default values. 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.

- 7. Optionally, clear the check box for **Auto start Virtual Machine** if you need to reprovision the VM before starting it for the first time.
- 8. Click **Create VM**. When the process is complete, the wizard closes automatically.
- If applicable, use the Reprovision Virtual Machine wizard to allocate additional resources to the VM, as described in Reprovisioning Virtual Machine Resources. Also, if you want to allocate additional space in each volume container for snapshots, see Expanding a Volume Container on the everRun System.

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

- 10. Click Console to open the console of the VM and log on to the guest operating system.
- 11. If necessary, update the network settings in the guest operating system.
Related Topics

Managing Snapshots Creating and Migrating Virtual Machines Managing the Operation of a Virtual Machine

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) or to a USB device. 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. For additional VM migration/export methods, see Creating and Migrating Virtual Machines.)

Prepare for exporting a snapshot by inserting a USB device or by creating a network-mounted folder to store an exported VM in your environment. If you are using a USB device, insert it into the primary node. If you are using a folder, 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. Then mount the folder or USB device 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 as standard Open Virtualization Format (OVF) and Virtual Hard Disk (VHD) files.

Notes:

Ö

- 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.
- 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 ever-Run System.
- The time required for the export depends on the size and number of volumes in the source VM as well as network bandwidth. For example, transferring a VM with one 20 GB boot disk over a 1Gb network may take about 30 minutes.
- If you will continue to use the source VM after 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.

Prerequisites:

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

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

To export a snapshot

- 1. Log on to the everRun system with the everRun Availability Console.
- 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.
- 3. If you have not already done so, create a snapshot, as described in Creating a Snapshot.
- 4. On the **Snapshots** page, select the snapshot to export.

Snapshots are usually in a normal state (\checkmark) the **State** column. If a snapshot is broken (\thickapprox), it may indicate that one or more volumes in the snapshot are unavailable for export. You can check volume availability in step 10

5. Click **Export** to open the export wizard.

- 6. Select one of the following:
 - Mount device via Windows Share (CIFS/SMB)

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

Mount device via NFS

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

Mount USB

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

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

If the mount succeeds, the repository appears under **Device URL** and the **Export VM** button becomes active; otherwise, an alert appears.

- 9. For All Captured Data Volumes are available for Export from node*n*, select the volumes to include. (The boot volume is required.)
- 10. Click Export Snapshot to export the VM.

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

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

The 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.

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

Troubleshooting

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

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

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

Related Topics

Managing Snapshots

Creating and Migrating Virtual Machines

Managing the Operation of a Virtual Machine

Removing a Snapshot

Remove a snapshot to permanently delete it from the everRun system. You can remove a snapshot from the **Virtual Machines** page or from the **Snapshots** page.

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 stor

age group.

To remove a snapshot (Snapshots page)

- 1. On the **Snapshots** page, select a snapshot to remove.
- 2. In the bottom pane, click **Delete**.
- 3. A confirmation window appears. Click **Yes** to continue deleting the snapshot or click **No** to stop the deletion.

To remove a snapshot (Virtual Machines page)

- 1. In the top pane of the **Virtual Machines** page, select the VM that contains the snapshot you want to remove.
- 2. In the bottom pane, click the **Snapshots** tab.

- 3. Select the snapshot that you want to remove.
- 4. In the Action column, click Remove.
- A confirmation window appears. Click Yes to continue deleting the snapshot or click No to stop the deletion.

Related Topics

Managing Snapshots

Creating and Migrating Virtual Machines

Managing the Operation of a Virtual Machine

Advanced Topics (Virtual Machines)

The following topics describe procedures and information for advanced users:

- Assigning a Specific MAC Address to a Virtual Machine
- Selecting a Preferred PM for a Virtual Machine
- Forcing a VM to Boot
- Changing the Protection Level for a Virtual Machine (HA or FT)
- Configuring the Boot Sequence for Virtual Machines
- Resetting MTBF for a Failed Virtual Machine
- Locating a Dump File in a Virtual Machine
- Attaching a USB Device to a Virtual Machine

To manage the operation of a virtual machine, see Managing the Operation of a Virtual Machine.

Assigning a Specific MAC Address to a Virtual Machine

Assign a specific Media Access Control (MAC) address to a virtual machine (VM) if you need to override its default MAC address.

Warnings:

- 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).
- 2. Select a VM and click **Shutdown**.
- 3. When the VM has stopped, click Config to display the Reprovision Virtual Machine wizard.
- 4. Click **Next** on each wizard page until the **Networks** page is displayed. (If applicable, see Reprovisioning Virtual Machine Resources to configure additional VM resources.)
- 5. On the **Networks** page, locate the network to modify and make a note of the current MAC address in case you need to restore it.
- 6. Type the new address in the **MAC address** column, or leave the text area blank to allow the everRun software to automatically assign the MAC address.
- 7. Click Finish.

Related Topics

Advanced Topics (Virtual Machines) Managing Virtual Machine Resources Managing the Operation of a Virtual Machine

Selecting a Preferred PM for a Virtual Machine

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

ñ

Advanced Topics (Virtual Machines)

Managing the Operation of a Virtual Machine

Forcing a VM to Boot

I.

You can force a VM to boot using the **Force Boot** button on the VIRTUAL MACHINES page. However, the **Force Boot** button is active only when the everRun Availability Console reports that the partner node is powered off or otherwise unreachable. When you use **Force Boot** to bring a VM online, you manually by-pass the system's safety checks to protect data, so you must use **Force Boot** with extreme caution and with full understanding of the conditions and consequences of using it.

Caution: Before using **Force Boot**, read this entire topic and consult with your authorized Stratus service representative. The service representative can review your system, including the date of the last volume synchronization, and can then discuss with you the full impact of using **Force Boot**. Then, you can decide, with your service representative, whether or not to force a VM to boot.

When you force a VM online with **Force Boot**, you select a node (that is, the node that is reachable) on which to force the VM to boot. All data on that node is marked as valid, regardless of the actual condition of the data (for example, the data's state, the last synchronization, the condition of the volume, etc.).

During the **Force Boot** process, the VM's volumes are tagged with the date and time that the force-boot process was initiated. The VM's AX components (that is, the VM's AX pair) use the data on the VM's volumes and communicate the status of that data to determine which AX contains the up-to-date volume information. The **Force Boot** process overrides the built-in logic that protects a VM from running in a split-brain condition. If the AX pair cannot communicate, a split-brain condition occurs and damages data integrity (for information on the split-brain condition, see Creating a SplitSite Configuration).

Warnings: Do not use Force Boot in the following situations:

- One or more volumes is the target of an unfinished mirror copy on the node where you will perform **Force Boot**.
- A target of an unfinished mirror copy is not good and will not be available even with Force Boot.
- The volumes are not synchronized. The following two situations are examples:
 - Both of the VM's AXs must have access to all of the VM's data volumes.
 - On a system with multiple volumes, the VM needs both AXs to be running in order for the VM to have access to all of its volumes because each node has a greenchecked copy of a different volume, and the volume's mirror copy on the opposite node is not green-checked.
- Both nodes are required because multiple VMs are degraded, yet are green-checked on opposite nodes (for example, Node0 has a good boot volume but a bad data volume, while Node1 has a bad boot volume but a good data volume).
- The system is licensed for one node.

If you perform a **Force Boot** on a system with outdated volumes, contact your authorized Stratus service representative immediately. If both nodes are powered on and have started to synchronize data, the system uses data from the VM that you forced to boot, and you cannot recover the data on the node that was unreachable. In some circumstances, however, you might be able to recover data after you use **Force Boot** on a system with outdated volumes:

- If the unreachable node is still powered off, do not power it on.
- If the unreachable node was powered off before you clicked Force Boot, then the VM's AX on the powered-off node is preserved and you can reverse the Force Boot without data loss under the following conditions:

- The VM that you forced to boot does not have new data (that is, the VM has not been put in production).
- Before you forced the VM to boot, the VM's AX on the unreachable node did not exchange status with the AX of the VM that you will force to boot.
- The issue preventing the VM's AX on the unreachable node from booting is resolved.
- All VM data between the two nodes is accurately synchronized. The system has no VMs where, of each VM's two AX components, the data of the VM's AX on one node is in a different state from the data of the VM's AX on the other node.

If your system meets all of the conditions above, contact your authorized Stratus service representative to advise you on a recovery process.

If you have decided to force a VM to boot, be sure to prepare for it by performing the prerequisite procedures.

Prerequisites: Manually check all volumes to ensure that you can safely override them. For example, the volume state should be green-checked, and disk synchronization should be finished. Determine if both AX components of the VM can communicate and can allow the system processes to determine the state of each volume. To prevent a split-brain condition, you must ensure that the two AX components of the VM can communicate status and can determine which AX has good data and good boot volumes. Ensure that the system is licensed for two nodes. Contact your authorized Stratus service representative.

To force a VM to boot

After you have consulted with your authorized Stratus service representative, and you have decided to force a VM to boot, perform the following procedure. In the examples, node0 is offline, node1 is the primary, and VM-1 is stopped.

- 1. In the everRun Availability Console of a system licensed for two nodes, click Virtual Machines in the left panel.
- 2. Navigate to the Virtual Machines page.
- 3. On the **Virtual Machines** page, select the VM that is stopped and that you want to force to boot (for example, VM-1).

4. In the bottom panel, click the **Start** button.

The VM begins to boot. It continues booting until the time-out limit is reached, possibly as long as 5 minutes. When the time-out limit is reached, the **Force Boot** button becomes active.

5. To force the VM to boot, click **Force Boot**.

A warning appears, asking you if you are positive that you know which node has the most up-to-date VM data. The warning also tells you to be aware that data loss is possible. In addition, a message tells you the node on which you can force the VM to boot.



You must type the node (node0 or node1) as indicated in the message. The following message is an example:

Force Boot VM-1	
DO NOT PROCEED UNLESS YOU	J ARE POSITIVE YOU KNOW
WHICH NODE HAS YOUR MOST	UP TO DATE VM DATA. BE
AWARE THAT DATA LOSS IS PO	SSIBLE.
Only node1 can be force-booted.	
If you would like to boot the VM on no	de1, type node1 :
node1	
[OK button]	[Cancel button]

 Click OK to force the node (for example, node1) to boot. (Click Cancel to cancel the procedure.) As the force-boot process begins and continues, additional confirmation messages appear before the VM starts and the data is marked as valid to the system.

The VM begins to run. On the **Virtual Machines** page, the VM is listed with a warning because the node (for example, node0) is still offline.

Once the secondary node is brought back in to the system, all data synchronizes from the node running the VM. In this example, all data synchronizes from node1 to node0.

Related Topics

Advanced Topics (Virtual Machines)

Managing the Operation of a Virtual Machine

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 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 (HA or FT)

Advanced Topics (Virtual Machines)

Managing the Operation of a Virtual Machine

Configuring the Boot Sequence for Virtual Machines

Configure the boot sequence of virtual machines to set the order in which guest operating systems and applications are started on the 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
Арр	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)

Managing the Operation of a Virtual Machine

Resetting MTBF for a Failed Virtual Machine

Reset the mean time between failure (MTBF) counter for a virtual machine to attempt to restart a failed virtual machine.

If a virtual machine's guest OS crashes, 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:

0

- 1. The **Reset Device** button is displayed only if the virtual machine falls below its MBTF threshold.
- 2. The **Clear MTBF** button is displayed only if the system software supporting a VM on one physical machine falls below its MBTF threshold.

To reset the MTBF counter for a virtual machine

- 1. On the Virtual Machines page, select a virtual machine.
- 2. Click Reset Device.

If the system software supporting a VM on one physical machine fails too often, perform the steps below to reset its MTBF counter.

To reset the MTBF counter for a VM on one physical machine

- 1. On the Virtual Machines page, select a virtual machine.
- 2. Click Clear MTBF.

Related Topics

Advanced Topics (Virtual Machines)

Managing the Operation of a Virtual Machine

Creating a Diagnostic File

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)

Managing the Operation of a Virtual Machine

Creating a Diagnostic File

Attaching a USB Device to a Virtual Machine

Attach a USB device to a virtual machine (VM) in order to enable the VM to use the device. A USB device may be needed, for example, when a USB-based license is required to install an application in a guest operating system. When you no longer need the USB device, detach it.

(If you need to mount a USB device on the everRun system to use the device for exporting or importing VMs, see Mounting a USB Device or Network-mounted Folder on the everRun System.)

Caution:

L

When you attach a USB device to 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, detach and remove the USB device as soon as you finish using it.



To attach a USB device to a VM

1. Insert the USB device into the primary (active) node for the VM.

The Virtual Machines page displays the primary node for each VM as the Current PM. (This node may be different from the current primary node for the everRun system, as displayed on the Physical Machines page.)

Confirm that the system displays the USB device. Navigate to the **Physical Machines** page. Click the node into which you inserted the device, and in the lower pane, select the **USB Device** tab. The USB device you inserted should appear in the tab's display.

- 2. On the Virtual Machines page, select a VM.
- 3. In the bottom pane, click the CD Drives & USB Devices tab.
- 4. On the USB line of the CD Drives & USB Devices tab, select a USB device from the pull-down menu.
- 5. Click Attach a USB to attach the USB device to the VM.
- 6. A **Confirm** dialog box appear, asking if you are sure you want to attach the device and displaying a warning that the guest goes simplex while the USB device is in use. Click **Yes** to attach the device.

After the system attaches the USB device to the VM, the USB device name appears in the list of USB devices on the **CD Drives & USB Devices** tab for the VM.

To detach a USB device from a VM

- 1. On the Virtual Machines page, select the VM to which the USB device is attached.
- 2. In the bottom pane, click the CD Drives & USB Devices tab.
- 3. On the USB line of the CD Drives & USB Devices tab, click Detach USB device. If necessary, select the USB device from the pull-down menu.

4. A **Confirm** dialog box appear, asking if you are sure you want to detach the device. Click **Yes** to detach the device.

After the system detaches the USB device to the VM, the USB device name no longer appears in the list of USB devices on the **CD Drives & USB Devices** tab for the VM.

Related Topics

Managing Virtual Machines

Chapter 8: Maintaining Physical Machines

You can maintain physical machines (PMs), or nodes, 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.

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 for hot-swappable components such as network cables, fans, and power supplies
 - Adding or Replacing Components That Are Not Hot-Swappable for components such as CPUs and memory or any other component that is not hot-swappable.
 - Adding a New NIC 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.
- To upgrade both PMs in a running system, see Upgrading Both Physical Machines In a Running System.

For information on disks, see Logical Disks and Physical Disks.

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 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 Maintaining Physical Machines The everRun Availability Console Physical Machines and Virtual Machines The Physical Machines Page

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.

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.
- 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 Maintaining Physical Machines The everRun Availability Console Physical Machines and Virtual Machines The Physical Machines Page

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.

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.
- 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 (¹ 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

Π

Maintenance Mode

Maintaining Physical Machines

The everRun Availability Console

Physical Machines and Virtual Machines

The Physical Machines Page

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.

To add new NICs

٥

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.
- 5. Reconfigure and start any VMs that need to use the new networks. See Managing Virtual Machines.

Related Topics

Maintenance Mode Maintaining Physical Machines The Physical Machines Page The Virtual Machines Page Business and Management Network Requirements General Network Requirements and Configurations

Replacing Physical Machines, Motherboards, NICs, or RAID Controllers

You replace motherboards, NICs, RAID controllers, and a physical machine (PM), or node, 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, including a failure of the motherboard, CPU, mid-plane, or storage controller. (If you need to recover the system software on a failed PM instead of replacing the PM hardware, see Recovering a Failed Physical Machine.)

When you remove and replace a PM, the system completely erases all of the disks in the replacement PM in preparation for a full installation of the everRun system software. To install the software, you can allow the system to automatically boot the replacement node from a temporary Preboot Execution Environment (PXE) server on the primary PM. As long as each PM contains a full copy of the most recently installed software kit (as displayed on the **Upgrade Kits** page of the everRun Availability Console), either PM can initiate the replacement of its partner PM with PXE boot installation. If needed, you can also manually boot the replacement node from DVD/USB installation media.

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

If you replace a PM or a component, use vendor instructions, but first read Physical Machine Hardware Maintenance Restrictions.

Caution: The replacement procedure deletes any software installed in the host operating system of the PM and all PM configuration information entered before the replacement. After you complete this procedure, you must manually re-install all of your host-level software and reconfigure the PM to match your original settings. **Caution**: To prevent data loss, if the system log indicates that manual intervention is necessary to assemble a disk mirror, contact your authorized Stratus service representative for assistance. You may lose valuable data if you force a resynchronization and overwrite the most recent disk in the mirror.

Prerequisite: If you want to use DVD or USB media to install the system software on the replacement PM, obtain the installation software for the release that the PM has been running by using one of the following methods:

- Create a bootable USB medium on the Upgrade Kits page, as described in Creating a USB Medium with System Software.
- Download an install ISO from your authorized Stratus service representative.
- Extract an install ISO into the current working directory from the most recently installed 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

If you download or extract an install ISO, save it or burn it to a DVD or USB medium. See Obtaining everRun Software.

Prerequisites: If you are replacing a PM, prepare the new PM:

- 1. Configure networks. See Network Architecture.
- 2. Configure storage. See Storage Requirements.
- 3. Connect power. See Connecting Power.
- Configure the firmware (BIOS or UEFI). See Configuring Settings in the Firmware Setup Utility.

Note:

 \mathcal{O}

Π

.

You must reactivate the product license for the everRun system after replacing a PM.

To remove and replace a failed PM or component (with PXE boot installation)

Use the following procedure to replace a failed PM, motherboard, NIC, or RAID controller and reinstall the system software by using PXE boot installation from the software kit on the primary PM.

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

Caution: Selecting **PXE PM Replace - Initialize All Disks** deletes all data on the replacement PM.

- 5. Select one of the following PXE Settings:
 - Only respond to PXE requests from the current partner node.

Waits for a PXE boot request from the MAC address of the current partner node. Select this option if you are recovering the existing PM by completely wiping and reinstalling it (with no hardware changes). This process deletes all data on the PM, but restores its current network configuration.

Only respond to PXE requests from the following MAC address.

Waits for a PXE boot request from the MAC address that you specify. Select this option if you are replacing the PM with a new PM, or replacing network adapters in the existing PM. Enter the MAC address of the specific network adapter that will initiate PXE boot.

Accept PXE requests from any system on priv0.

Waits for a PXE boot request from priv0, the private network that connects the two ever-Run nodes. Select this option if you are replacing the PM with a new PM, or replacing network adapters in the existing PM, but you do not know the MAC address for the new PM.

6. If prompted, under Assumed Network Settings, select one of the following options:

- Use below settings—The PM uses the displayed network settings. No user interaction is needed during the software installation process.
- Ask during install—The PM prompts for network settings. When the software installation begins, you must be present at the console of the replacement PM to enter the settings.

- Click Continue to begin the replacement process. The system shuts down and powers off the PM.
- 8. After the PM is powered off, install the replacement PM or component, if applicable:
 - a. 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. Connect a monitor and keyboard.
 - b. Reconnect the network cables to their original ports. Check that Ethernet cables are connected from the replacement 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.
- 9. Manually power on the replacement PM. As the PM powers on, enter the firmware (BIOS or UEFI) setup utility, and enable PXE boot (boot from network). If you previously selected Only respond to PXE requests from the following MAC address, enable PXE boot on the NIC associated with that MAC address; otherwise, verify that PXE boot is enabled on the priv0 NIC. Save the setting and restart the system.
- 10. The replacement process continues, as follows:
 - The replacement PM begins to boot from a PXE server that temporarily runs on the primary node.
 - The system automatically deletes all of the data on disks in the replacement PM.
 - The replacement PM reboots again and automatically starts the system software installation, which runs from a copy of the installation kit on the primary node.

If you previously selected **Ask during install** to specify the network settings of the replacement PM during the installation, monitor the installation process and respond to prompts at the physical console of the replacement PM; otherwise, skip to step 16.

11. 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:

- 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.
- 12. 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.

i

П

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.

- 13. 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.
- 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 redisplays until you enter valid information.

- 15. At this point, the software installation continues without additional prompts.
- 16. When the software installation is complete, the replacement PM reboots from the newly installed system software.



0

٥

ñ

Note: After the system software installation, the replacement PM may take up to 20 minutes to join the system and appear in the everRun Availability Console.

- 17. As the replacement PM joins the system, you can view its activity on the Physical Machines page of the everRun Availability Console. The Activity column displays the PM as (in Maintenance) after the recovery is complete.
- 18. Assign logical disks from the replacement PM to storage groups on the everRun system, as described in Assigning a Logical Disk to a Storage Group.

Notes:

- When the replacement 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 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 replacement PM to the same storage groups; otherwise, the everRun system cannot fully synchronize.
- To activate the replacement PM, re-activate the product license for the everRun system. On the Preferences page, click Product License, expand License Check and Activation, and click Check License Now to automatically activate the license (as described in Managing the Product License).

Note: The new PM cannot exit maintenance mode and run VMs until the everRun license is re-activated.

ñ

- 20. If applicable, manually reinstall applications and any other host-level software, and reconfigure the replacement PM to match your original settings.
- 21. 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.

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

To remove and replace a failed PM or component (with DVD/USB installation)

Use the following procedure to replace a failed PM, motherboard, NIC, or RAID controller and reinstall the system software by using a DVD or USB medium.

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



- 5. Click **Continue** to begin the replacement process. The system shuts down the PM in preparation for the system software reinstallation.
- 6. After the PM is powered off, install the replacement PM or component, if applicable:
 - a. 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. Connect a monitor and keyboard.
 - b. Reconnect the network cables to their original ports. Check that Ethernet cables are connected from the replacement PM (or new NIC) to the network or directly to the running (primary) PM, if the two everRun system PMs are in close proximity. One Ethernet cable should connect from the first embedded port on the new PM or from a NIC port if the new PM does not have an embedded port.

- 7. Insert the bootable media or mount the ISO image on the replacement PM, and then manually power on the PM.
- As the replacement PM powers on, enter the firmware (BIOS or UEFI) setup utility and set the Optical Drive or USB media as the first boot device.
- 9. Monitor the installation process at the physical console of the replacement PM.

Note: If necessary, see Installing Software on the Second PM for reference.Although that topic refers to the "second PM," in this instance, it applies to the replacement PM.

- 10. At the **Welcome** screen, use the arrow keys to select the country keyboard map for the installation.
- 11. At the **Install or Recovery** screen, select **Replace PM**, **Join system: Initialize Data** and press **Enter**.

Caution: Selecting Replace PM, Join system: Initialize data deletes all data on the replacement PM.

12. 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:

0

П

- 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.
- 13. The Select interface for managing the system (ibiz0) screen sets the physical interface to

П

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.

- 14. 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.
- 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

0

- Default gateway address
- Domain name server address

See your network administrator for this information.

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

- 16. At this point, the software installation continues without additional prompts.
- 17. When the software installation is complete, the replacement PM reboots from the newly installed system software.

Note: After the system software installation, the replacement PM may take up to 20 minutes to join the system and appear in the everRun Availability Console.

- As the replacement PM joins the system, you can view its activity on the Physical Machines page of the everRun Availability Console. The Activity column displays the PM as (in Maintenance) after the recovery is complete.
- 19. Assign logical disks from the replacement PM to storage groups on the everRun system, as described in Assigning a Logical Disk to a Storage Group.

	Notes	S.
i	•	When the replacement PM joins the everRun system, the system auto- matically adds the secondary everRun system disk to the Initial Storage Group; however, the system does not assign any other logical disks from the 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 replacement PM to the same storage groups; otherwise, the everRun sys- tem cannot fully synchronize.

20. To activate the replacement PM, re-activate the product license for the everRun system. On the **Preferences** page, click **Product License**, expand **License Check and Activation**, and click **Check License Now** to automatically activate the license (as described in Managing the **Product License**).



Note: The new PM cannot exit maintenance mode and run VMs until the everRun license is re-activated.

- 21. If applicable, manually reinstall applications and any other host-level software, and reconfigure the replacement PM to match your original settings.
- 22. 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.

Related Topics

Maintenance Mode

Maintaining Physical Machines

The everRun Availability Console

Physical Machines and Virtual Machines

The Physical Machines Page

Upgrading Both Physical Machines In a Running System

Prerequisite: Before you upgrade to new physical machines, see Physical Machine Hardware Maintenance Restrictions.

To upgrade to new physical machines

- Upgrade everRun software if required to support the new PM. See the appropriate everRun Release Notes and The Upgrade Kits Page.
- 2. Upgrade the first PM; see Replacing Physical Machines, Motherboards, NICs, or RAID Controllers.
- 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.

Related Topics

Maintenance Mode Maintaining Physical Machines The everRun Availability Console Physical Machines and Virtual Machines The Physical Machines Page

Part 2: Supporting Documents

See the following support documents for release information, and reference and troubleshooting information.

- everRun Release 7.5.1.0 Release Notes
- everRun Command Line Interface Reference
- System Reference Information
- SNMP

Chapter 9: everRun Release 7.5.1.0 Release Notes

These Release Notes are for everRun Release 7.5.1.0 (updated at 9:22 PM on 4/29/2021). See the following sections:

- Accessing Stratus Knowledge Base Articles
- Bug Fixes
- New Features and Enhancements
- Important Considerations
- Known Issues
- Documentation Updates
- Getting Help

Accessing Stratus Knowledge Base Articles

The **Stratus Customer Service Portal** provides a searchable **Knowledge Base** with technical articles about all Stratus products, including everRun. In some cases, the Release Notes directly reference these Know-ledge 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 7.5.1.0

EV-46765: Nodes may fail to return to service after power loss.

EV-47079: A fresh install of everRun Release 7.5.0.5 can leave a second logical disk stuck as foreign on one node if 4K data drives are present.

EV-47788: The eAC may become inaccessible after system power loss.

EV-47803: A spine restart may cause a false "Node has Rebooted" alert.

EV-47824: The fatal alert "A sensor on node0 (or 1) is indicating a fatal problem" can be seen when it should not be.

EV-47853: Upgrading to everRun Release 7.5.0.5 on a system with 4K data drives can leave logical disks in an unusable state that if repaired can lead to data loss.

EV-47858: Insulate volume replication status and epoch values against node power loss.

EV-47865: The heartbeat configuration file /etc/ha.d/ha.cf can be truncated on node power loss resulting in heartbeat and spine failing to start and the node failing to join the cluster.

New Features and Enhancements

Major new features and enhancements are listed below under the release in which they became available.
New in everRun Release 7.5.1.0

everRun Release 7.5.1.0 is for bug fixes. See Bug Fixes for information.

New in everRun Release 7.5.0.5

everRun Release 7.5.0.5 includes the following new features:

- OEM Support
 - Enhanced customization capabilities
 - International keyboards during installation and in host OS
- Performance
 - Upgraded host OS (CentOS 7.5)
 - Fault Tolerant (FT) performance improvements
 - Networking performance improvements
 - Configurable Maximum Transmission Unit (MTU) for networks in everRun Availability Console
 - Upgraded Quick EMUlator (QEMU) hypervisor
 - Support for QEMU Copy on Write (QCOW3, or QCOW2v3) disk images
 - Delta-copy disk synchronization during node recovery
 - Support for 25 Gb A-Links
- everRun Availability Console
 - New VM Manager user role
 - Enhanced performance statistics on System page
 - Improved usability and serviceability
- Installation and Upgrade
 - Administrator-controlled upgrades (full control of upgrade steps)
 - Reasons for Qualify failures reported in everRun Availability Console
 - Access to Back button during installation

- Flexibility
 - New guest OSes (see Compatible Guest Operating Systems)
 - Recover failed node from software kit on healthy node (enhanced PXE recovery)
 - Mount USB storage devices in VM
 - Assign custom network names
 - Rename VM while creating VM from snapshot
 - Support for NVM Express (NVMe) disk drives as data disks
- VM Management
 - Configure Auto Start Mode for VMs
 - Browse network shares for OVF files in Import/Restore VM Wizard
 - Migrate VMs from VMware ESXi with the P2V client (virt-p2v)
 - Improved capacity planning for VM resources
 - View per-volume progress while exporting VMs
 - Mount external USB drive to export VMs
 - Export Snapshots to USB drive
 - Enable/disable VM components
- Operations and Troubleshooting
 - SNMP updates
 - Set MAC address when creating VM
 - Cancel a snapshot

Important Considerations

Upgrading from everRun Release 7.4.x.x

You can upgrade to the current release from everRun Release 7.4.*x.x* by following the instructions in Upgrading everRun Software.

Caution: All PMs and VMs must be in good health before you upgrade 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 Release 7.4.*x*.*x* to Release 7.5.*x*.*x*, consider the following:

• External storage is not supported.

everRun Release 7.5.*x.x* or higher no longer supports external storage. If you need to upgrade a system that includes external storage, contact your authorized Stratus service representative for assistance.

• Eject any VCDs or USB media from the VMs before upgrading.

If VCD or USB media is still connected to the VMs, it prevents the system from migrating the VMs and putting the PMs into maintenance mode for the upgrade process.

• Update the Red Hat VirtIO drivers in any existing Windows-based VMs after upgrading the system.

After upgrading to Release 7.5.*x.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).

• Reboot any existing VMs to take advantage of performance improvements.

After upgrading to Release 7.5.x.x, reboot the VMs on the system at your earliest convenience to take advantage of the performance improvements in the upgraded QEMU hypervisor.

• Refresh web browser after upgrade to ensure up-to-date everRun Availability Console displays.

After an upgrade, various console displays may be incorrect due to files cached in the web browser. For example, the **Reprovision Virtual Machine** wizard may display incorrect values. To correct the display, reload it by pressing **Ctrl-F5**.

Upgrading from everRun Release 7.3.4.x or Earlier

L

You can upgrade to the current release from everRun Release 7.3.4.x by following the instructions in Upgrading everRun Software. If you are upgrading from an everRun release prior to Release 7.3.4.x, you must upgrade to Release 7.3.4.x first, and then upgrade to Release 7.5.x.x.

Caution: All PMs and VMs must be in good health before you upgrade 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.*x* or previous releases to everRun Release 7.5.*x*.*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.x* or higher no longer supports Nagios monitoring, which was a technology preview in everRun 7.3.*x.x*. Nagios add-ons are automatically removed during the upgrade process.

• External storage is not supported.

everRun Release 7.5.*x*.*x* or higher no longer supports external storage. If you need to upgrade a system that includes external storage, contact your authorized Stratus service representative for assistance.

• everRun Release 7.4.x.x or higher requires more disk space than previous releases.

everRun Release 7.4.*x.x* or higher requires more disk space than previous releases to accommodate changes for the host OS. The amount of space required to upgrade from Release 7.3.4.*x* to 7.5.*x.x* can be up to 10 GiB per logical system disk. To verify that your Release 7.3.4.*x* system meets the upgrade requirements, upload the Release 7.5.*x.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.*x* or higher.) For more information about using the **Qualify** button and upgrading your system, see Upgrading everRun Software.

• Eject any VCDs or USB media from the VMs before upgrading.

If VCD or USB media is still connected to the VMs, it prevents the system from migrating the VMs and putting the PMs into maintenance mode for the upgrade process.

- Update the Red Hat VirtIO drivers in any existing Windows-based VMs after upgrading the system.
 After upgrading to Release 7.5.*x.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).
- Any manually added RPMs or customized host OS configuration settings are lost.

When you upgrade from Release 7.3.4.*x*, the system 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, SSH keys) 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.x* or higher 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.x* or higher.

Re-activating the everRun Product License After Replacing a PM

When you replace a PM in an everRun system, you must re-activate the product license for the system to register the new PM with the system. The new PM cannot exit maintenance mode and run VMs until the ever-Run license is re-activated.

After replacing the PM, when the new PM is **running (in Maintenance)** on the **Physical Machines** page, open the **Preferences** page, click **Product License**, and expand **License Check and Activation**. Click **Check License Now** to automatically activate the license (as described in Managing the Product License). When the license is activated, click **Finalize** to exit maintenance mode.

Other Important Considerations

For other important considerations about everRun systems, see Important Physical Machine and Virtual Machine Considerations.

Known Issues

P2V/V2V Migration May Fail when Migrating a Windows-based VM

If you use the P2V client (virt-p2v) to migrate a PM or VM running Windows, the migration may fail with the message, "error: inspection could not detect the source guest (or physical machine)." For more information about this issue, see KB-9401.

VM Network Remains Offline After Using P2V/V2V to Migrate from a VMware Source

If you use the P2V client (virt-p2v) to migrate a PM or VM from a source that is running VMware, the network driver might fail to install properly during the migration process, which leaves the VM network offline. If this happens, manually install the network driver in Device Manager. For instructions, see the **Troubleshooting** section of Migrating a Physical Machine or Virtual Machine to a System.

VM Network Remains Offline After Importing a Windows-based VM

If you import a Windows-based VM with the **Import/Restore Virtual Machine Wizard**, the VirtIO network driver may fail to install properly the first time you start the VM, which leaves the VM network offline. If this happens, manually install the network driver in Device Manager. For instructions, see the **Troubleshooting** section of Importing an OVF File.

Ensure that Both Nodes Are Running Before Removing a VM

To avoid downtime, before removing a VM, always ensure that both nodes are in the **running** state and that neither node is in maintenance mode or in the process of synchronizing.

If you remove a VM while one node is shut down in maintenance mode, the operation may stall with the error, "There was a problem deleting *vm-name*." To correct the problem, you must restart the only running node.

Invalid VM Names Prevent Some VM Operations from Executing

The 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 must start with a character within the range [a-z], [A-Z], or [0-9].
- A VM name cannot start with the prefixes <code>Zombie-ormigrating-</code>, which are reserved strings.
- A VM name can contain characters only within the range [a-z], [A-Z], or [0-9]; or _ (underline), (minus), or . (dot).
- A VM name must be 1 to 64 characters in length.

Guest OS May Fail to Display a VCD When Inserted

After upgrading an everRun system from Release 7.3.4.*x* or earlier to Release 7.5.*x*.*x*, if you insert a virtual CD (VCD) into a VM as described in Inserting a Virtual CD, the guest operating system may fail to display or

mount the VCD.

This issue occurs because the guest OS is still running the QEMU hypervisor from the older version of ever-Run. To correct the issue, restart the guest operating system, which automatically loads the updated QEMU hypervisor. At your convenience, ensure that you restart the guest OS of each VM on your system to run the updated QEMU hypervisor.

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 **VNC: Unable to connect, retrying in** *n* **seconds** may appear. To enable the VNC connection, click the https link to the VNC console page in the upper right-hand corner of the masthead, and continue with the appropriate procedure below:

- In Internet Explorer, the Security Alert wizard appears:
 - a. Click Continue to this website (not recommended).
 - b. Click OK.
- In FireFox, the Your connection is not secure window appears:
 - a. Click Advanced. A message about an invalid security certificate appears.
 - b. Click Add Exception. The Add Security Exception dialog box appears with the console's location in Location.
 - c. Click Confirm Security Exception.

The VNC console appears.

Cannot Mount USB Media with exFAT File System

The everRun software on everRun systems does not support the exFAT File system. If you need to mount a USB medium to import or export a VM, as described in Mounting a USB Device or Network-mounted Folder on the everRun System, format the device with NTFS. (By default, most USB media are formatted with the FAT file system, which has a limited file size of 4 GB that may be too small for most VMs.)

Recovering a Missing or Failed Logical Disk with the Repair Button May Be Slow

If you attempt to recover a missing or failed logical disk with the **Repair** button in the masthead of the ever-Run Availability Console, as described in Responding to a Failed Logical Disk, the system may be slow to repair the disk. Although the system successfully removes the failed logical disk from its storage group, it may be slow to migrate the data from the failed disk to other disks in the storage group. The **Alerts** page may continue to report that the logical disk is not present, that volumes have failed, and that storage is not fault tolerant. Also, the **Volumes** page may continue to show volumes in the broken (**X**) state. If this condition persists, contact your authorized Stratus service representative for assistance.

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 QCOW3 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 QCOW3 (QCOW2v3) format. The conversion is expected because snapshots require QCOW3 format, but consider that changing the volume format may have performance implications for certain VM loads. You cannot convert a volume from QCOW3 back to RAW format. 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.

Reboot Required when Changing Node IP Address or Netmask Network Settings

When you change the IP address or netmask settings of an everRun node as described in Configuring IP Settings, you must reboot the node or the system for the changes to take effect. If you do not reboot the node or the system, both the new and old settings will remain in effect, which may cause routing or connection issues.

Changing the MTU of Business Networks Causes VM Migration and Possibly Node Failover

If you change the maximum transmission unit (MTU) setting of either business network (network0 or network1) as described in Setting the MTU, the system automatically migrates the VMs from one node to the other. If you change the MTU for network0 specifically, the system also automatically fails over from the primary node to the secondary node.

To avoid this issue, avoid changing the MTU setting of the business networks, or change the MTU only during a planned maintenance period.

Documentation Updates

Updates in everRun Release 7.5.0.5 Documentation

The following changes were made late in the release:

- Migrating a Physical Machine or Virtual Machine to a System lists additional guest operating systems supported for P2V migration, and describes how to properly disable hibernation and fast startup settings before migrating Windows-based systems.
- Business and Management Network Requirements warns against using the ifdown command to temporarily bring down a VM's business network connection.

Getting Help

If you have a technical question about everRun systems, you can find the latest technical information and online documentation at the **Downloads** page at https://www.stratus.com/services-sup-port/downloads/?tab=everrun. You can also search the **Knowledge Base** in the **Stratus Customer Service Portal** at https://support.stratus.com.

If you cannot resolve your questions with these online resources, and the system is covered by a service agreement, contact your authorized Stratus service representative. For information, see the **everRun Support** page at https://www.stratus.com/services-support/customer-support/?tab=everrun.

10

Chapter 10: 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
- AVCLI Command Descriptions

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
- Installing the Linux Client
- Installing the Windows Client

The following topics explain how to use the AVCLI command interface:

- Using AVCLI
- Executing a Command
- Using AVCLI Help

The following topics are helpful for programmers using the AVCLI command interface:

- AVCLI Error Status
- XML Encapsulated Errors

- Error Checking
- Asynchronous Command Delay
- Output Formatting
- AVCLI Exceptions

Related Topics

AVCLI Command Descriptions

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:
 - Go to the Downloads page at https://www.stratus.com/services-support/downloads/?tab=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:
 - Go to the Downloads page at https://www.stratus.com/services-support/downloads/?tab=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:

П

- 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 hostname, 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 hostname, username, and password, as described in Executing a Command.

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.

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::SimpleLibrary. 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:0406
```

```
-> 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

ň

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>
<node>
<iavance>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>

ň

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.throwNonE
xistentResource(ResourceDisambiguateServiceProvider.java:56)
at
com.avance.yak.cli.ResourceDisambiguateServiceProvider.getResour
ceId(ResourceDisambiguateServiceProvider.java:81)
at
com.avance.yak.cli.Command.findResourceId(Command.java:80)
at
com.avance.yak.cli.CommandWithUnparsedAmbiguousResourcesInvokeEa
ch.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.throwNonE
xistentResource(ResourceDisambiguateServiceProvider.java:56)
at
com.avance.yak.cli.ResourceDisambiguateServiceProvider.getResour
ceId(ResourceDisambiguateServiceProvider.java:81)
at
com.avance.yak.cli.Command.findResourceId(Command.java:80)
at
com.avance.yak.cli.CommandWithUnparsedAmbiguousResourcesInvokeEa
ch.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 short-0 cut as described in Executing a Command. Help help **Basic System Information** audit-export audit-info unit-avoid-bad-node unit-change-ip unit-configure unit-eula-accept unit-eula-reset unit-info unit-shutdown unit-shutdown-cancel unit-shutdown-state unit-synced **System Configuration** callhome-disable callhome-enable callhome-info datetime-config dialin-disable

dialin-enable

dialin-info

ealert-config

ealert-disable

ealert-enable

ealert-info

license-info

license-install

ntp-config

ntp-disable

proxy-config

proxy-disable

proxy-enable

proxy-info

snmp-config

snmp-disable

snmp-info

timezone-config

timezone-info

System User Management

ad-disable

ad-enable

ad-info

ad-join

ad-remove

local-group-add

local-group-delete

local-group-edit

local-group-info

local-user-add

local-user-delete

local-user-edit

local-user-info

owner-config

owner-info

Managing Physical Machines

node-add

node-cancel

node-delete

node-info

node-reboot

node-recover

node-shutdown

node-workoff

node-workon

pm-clear-mtbf

Managing Alerts

alert-delete

alert-info

Diagnostic Files

diagnostic-create

diagnostic-delete

diagnostic-extract

diagnostic-fetch

diagnostic-info

Kit Information

kit-add

- kit-controlled-upgrade-continue
- kit-controlled-upgrade-disable
- kit-controlled-upgrade-enable
- kit-controlled-upgrade-info
- kit-delete
- kit-info
- kit-qualify
- kit-upgrade
- kit-upgrade-cancel

Network/Storage Information

- disk-move-to-group
- image-container-info
- image-container-resize
- network-change-mtu
- network-change-role
- network-info
- node-config-prp
- node-delete-prp
- removable-disk-info
- storage-group-create
- storage-group-delete
- storage-group-info
- storage-info
- volume-info
- volume-resize

Creating Virtual CD/DVDs

- media-create
- media-delete

media-eject

media-import

media-info

media-insert

Managing Virtual Machines

localvm-clear-mtbf

ova-info

ovf-info

vm-attach-usb-storage

vm-ax-disable

vm-ax-enable

vm-boot-attributes

vm-cd-boot

vm-copy

vm-create

vm-create-from-snapshot

vm-delete

vm-device-config-info

vm-export

vm-import

vm-info

vm-media-insert-disable

vm-media-insert-enable

vm-network-disable

vm-network-enable

vm-poweroff

vm-poweron

vm-reprovision

vm-restore

vm-shutdown

vm-snapshot-create

vm-snapshot-create-disable

vm-snapshot-create-enable

vm-snapshot-delete

vm-snapshot-export

vm-snapshot-info

vm-unlock

vm-usb-attach-disable

vm-usb-attach-enable

vm-volume-disable

vm-volume-enable

Related Topics

AVCLI Command Overview

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

username <i>name</i>	The user with rights to join to the specified domain.
password password	The password of the user with rights to join to the spe- cified domain. If you do not give the password, you are automatically prompted for it.
secure-mode true false	Enable (true, the default) or disable (false) secure mode.
domain	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

username <i>name</i>	The user with rights to remove the everRun system from the specified domain.
password password	The password of the user with rights to remove the ever- Run system from the specified domain. If you do not give the password, you are automatically prompted for it.
secure-mode true false	Enable (true, the default) or disable (false) secure mode.
domain	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

alerts	One or more alerts to be deleted.
purge	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

alerts 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 ${\tt audit-info}$ command displays either the last 50 audit logs or the specified number of audit logs.

Options

	number-of-audit-logs	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

date	The date, formatted as YYYY-MM-DD.
time	The time, formatted as <i>HH</i> : <i>MM</i> : <i>SS</i> , in a 24-hour format.
timezone	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.

minimal	The smallest diagnostic (approximately 2 to 10 MB).
medium	A medium diagnostic (approximately 10 MB).
full	A large diagnostic (approximately 60 MB).

diagnostic-delete

Usage

avcli diagnostic-delete diagnostics...

Description

The diagnostic-delete command deletes the specified diagnostic files.

diagnostics	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.

diagnostics	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, full, dumps.
- *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.

\sim			
()	nti	n	٦C
\mathbf{v}	νu	υı	13

diagnostics	One or more diagnostic files to be downloaded.
file name	The name of the file written to the current directory. This option is valid if only one diagnostic is downloaded.
extract	Extract the downloaded diagnostic file(s.

)

Examples

- \$ avcli diagnostic-fetch buggrab:o10
- \$ avcli diagnostic-fetch --file buggrab.zip buggrab:010

\$ avcli diagnostic-fetch buggrab:o10 buggrab:o11 buggrab:o12

diagnostic-info

Usage

avcli diagnostic-info diagnostics...

Т

Description

The diagnostic-info command displays information about all diagnostics or only about those specified.

diagnostics	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 ${\tt dialin-info}$ 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.

disk	One or more disks to be moved.
storage-group	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.

ssl	Use SSL when communicating with the SMTP server. You can specify this option only if tls is not specified.
tls	Use TLS when communicating with the SMTP server(s). You can specify this option only if ssl is not specified.
port <i>polt</i>	Use this port number to connect to the SMTP server.
sender <i>sender</i>	Send email from this sender (name or IP address).
username <i>username</i>	Use this name for authentication on the specified SMTP host.
password password	Use this password for authentication on the specified SMTP host.
host <i>host</i>	The DNS or IP address of the SMTP server.
recipients	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.mydomain.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.mydomain.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 ${\tt help}$ command provides help about a specific command or lists all AVCLI commands.

Options

|--|

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

	The name of the image container. If you do not provide this argu-
image-container	ment, the command displays information about all image con-
	tainers.

Examples

```
$ avcli image-container-info
```

image-container:

- -> name : root
- -> id : imagecontainer:058
- -> 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:036
- -> hasFileSystem : false
- -> isLocal : true
- -> size : 2,151,677,952
- -> size-used : 2,151,677,952
- -> storage-group : none

image-container:

- -> name : swap
- -> id : imagecontainer:066
- -> 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:077
- -> 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

new-size <i>size</i>	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).
image-container	The name of the image container.

Examples

\$ avcli image-container-resize --new-size 40G boot-chom1

kit-add

Usage

```
avcli kit-add kit_path...
```

Description

The kit-add adds (that is, uploads) one or more upgrade kits to the Upgrade Kits page.

Options

kit_path	One or more upgrade kits to be added. This value is a path to a file.

Examples

\$ avcli kit-add everRun upgrade-7.5.0.0 0-129.kit

kit-controlled-upgrade-continue

Usage

avcli kit-controlled-upgrade-continue

Description

The kit-controlled-upgrade-continue command enables a controlled upgrade to continue with the next operation, after a pause in the upgrade process.

kit-controlled-upgrade-disable

Usage

avcli kit-controlled-upgrade-disable

Description

The kit-controlled-upgrade-disable command disables the system's ability to perform a controlled upgrade. After issuing this command, the kit-upgrade command performs a normal upgrade, not a controlled upgrade.

kit-controlled-upgrade-enable

Usage

```
avcli kit-controlled-upgrade-enable
```

Description

The kit-controlled-upgrade-enable command enables controlled upgrades on the system. After issuing this command, the kit-upgrade command performs a controlled upgrade..

In a normal upgrade, the console is locked during the entire upgrade. In a controlled upgrade, the upgrade process pauses at every transition into and out of maintenance mode, with a pop-up window displaying a message indicating that a controlled upgrade is paused and control buttons that allow you to select the next available action.

Controlling an upgrade can be useful for verifying or reconfiguring third-party tools or other system services that are not managed by the everRun system.

kit-controlled-upgrade-info

Usage

avcli kit-controlled-upgrade-info

Description

The kit-controlled-upgrade-info command displays information about the controlled upgrade.

Examples

The following is the command with sample output:

```
[root@node0 ~]# avcli kit-controlled-upgrade-info
Feature enabled : No
Toggle allowed : Yes
State : IDLE
Current action : None required.
```

In the output, the state and current action fields indicate the next expected action, which is typically to put a node into or out of maintenance mode. Issue the kit-controlled-upgrade-continue command to perform the next action.

kit-delete

Usage

avcli kit-delete kit_id

Description

The kit-delete command deletes the specified kit(s).

Options

kit_id	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 command description.

Example

```
kit-delete kit:024
```

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

Examples

You can issue the kit-info command to obtain the *kit-id* value for commands such as kit-upgrade, kit-qualify, and kit-delete. 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:024:

```
[root@node0 ~]# avcli kit-info
-> name : unspecified
-> id : kit:o24
-> description : unspecified
-> version : 7.5.0-127
-> locked : false
```

kit-qualify

Usage

```
avcli kit-qualify kit_id
```

Description

The kit-qualify command qualifies the specified upgrade kit file. If the qualification succeeds, the kit can successfully upgrade the system. 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 the disk has insufficient space to complete the upgrade, the file contains the message Insufficient free space and reports 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.

Options

kit_id The upgrade kit to be qualified. The value is the kit ID.

For information on obtaining the value for *kit_id*, see the kit-info command description.

Examples

```
kit-qualify kit:024
```

kit-upgrade

Usage

```
avcli kit-upgrade kit_id
```

Description

The kit-upgrade command begins an upgrade using the specified kit. After you issue the command, the prompt returns when the upgrade begins.

Options

kit_id The kit to use for the upgrade. The value is the kit ID.

For information on obtaining the value for *kit_id*, see the kit-info command description.

Examples

```
kit-upgrade kit:024
```
kit-upgrade-cancel

Usage

```
avcli kit-upgrade-cancel kit_id
```

Description

The kit-upgrade-cancel command cancels a kit upgrade. This command is effective only if you issue it before the first node is placed into maintenance mode during an upgrade.

Options

kit_id The kit upgrade to cancel. The value is the kit ID.

For information on obtaining the value for *kit_id*, see the kit-info command description.

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

 $The \verb+license-install+ command installs the specified license file.$

Options

	license-file	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

name <i>name</i>	Local group name.
permissions <i>permission-type</i>	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



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

name <i>name</i>	New local group name.
permissions <i>permission-type</i>	Local group permissions, in the form of a comma-separated list.
group-name-or-sid	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.

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.

username <i>name</i>	everRun local user name.
password password	Boolean flag that indicates whether the user should be prompted for a new password.
new-password password	Specify the password as a command-line option instead of being prompted in the same way aspassword.
realname <i>name</i>	The user's real name.
email <i>address</i>	The user's email address.
local-groups <i>groups</i>	Local groups for the user to join, in the form of a comma-separated list.
permissions <i>permission-types</i>	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

Page 386 of 539

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



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.

username <i>name</i>	The user name to assign.
password password	Boolean flag that indicates whether the user should be prompted for a new password.
new-password password	Specify the password as a command-line option instead of being prompted in the same way as - -password.
realname <i>name</i>	The user's real name.
email <i>address</i>	The user's email address.
local-groups <i>groups</i>	Local groups for the user to join, in the form of a comma-separated list.
permissions <i>permission-types</i>	Local user permissions, in the form of a comma- separated list.
group-name-or-sid	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.



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

storage-group group	The storage volume to carve from. If you do not specify this option, the storage group with the most free space is automatically selected.
name <i>name</i>	The name of the carved volume. If you do not specify this option, the name is determined from the URL.
url	The URL where the ISO file is located.
wait	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

 $The {\tt media-delete} \ command \ deletes \ the \ specified \ media.$



media-eject

Usage

avcli media-eject [--cdrom name] [vm...]

Description

The media-eject command ejects media from the specified virtual machines.

cdrom <i>name</i>	The CD-ROM device to eject. This value is optional if the VM has only a single CD-ROM device.
vm	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

storage-group group	The storage volume to carve from. If you do not specify this option, the shared storage with the most free space is automatically selected.
name <i>name</i>	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.
throttle	 Slow down the import/export operation. Valid values are: none: Do not use throttling. This is the default value. low: Slow down by about 25%. medium: Slow down by about 50%. high: Slow down by about 75%.
silent	Suppress output.
file	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.



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

L

iso name	The ISO image to insert.
cdrom <i>name</i>	The CD-ROM device to insert. This value is optional if the VM has only a single CD-ROM device.
vm	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

name	The name of the network
size	The MTU size. Valid values are 1280 - 65535 (1500 is the default).

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.

networks	One or more networks whose role is to be changed.
role	The new role. Specify either business or a-link.

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



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:
```

: sync_2003
: sharednetwork:o2334
: ft
: a-link
: 10 Gb/s
: 1500
: network0
: sharednetwork:064
: ft
: business
: 1 Gb/s

shared network: : sync_2004 -> name -> id : sharednetwork:o2333 -> fault-tolerant : ft -> role : a-link -> bandwidth : 10 Gb/s : 1500 -> mtu shared network: -> name : priv0 -> id : sharednetwork:065 -> 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.

wait	Wait for the command to complete.
-w	

node-cancel

Usage

avcli node-cancel pm

Description

The node-cancel command cancels a PM that is being imaged.



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

nicl adapter	The name of a physical adapter.
nic2 adapter	The name of a physical adapter.
node	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.

pm	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

name adapter	The name of the adapter to delete.
node	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.

	pm	The PM about which to display information.
--	----	--

node-reboot

Usage

avcli node-reboot [--wait] pm

Description

The node-reboot command reboots the specified PM.

wait -w	Wait for the command to complete.
pm	The PM to reboot.

node-recover

Usage

avcli node-recover [--wipe] pm [--wait]

Description

The node-recover command recovers the specified PM.

pm	The PM to recover.
wipe	Wipe the disks off the PM prior to recovery.
wait -w	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 or use the everRun Availability Console

Options

force -f	Override the shutdown warning.
wait -w	Wait for the command to complete.
pm	The PM (for example, node1) 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.

wait -w	Wait for the command to complete.
pm	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

pm	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

	servers	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 ${\tt ova-info}$ command displays information about the specified OVA files.

<i>filename</i> .ova	One or more OVA files.
----------------------	------------------------

ovf-info

Usage

avcli ovf-info filename.ovf...

Description

The ${\tt ovf-info}$ command displays information about the specified OVF files.

<i>filename.</i> ovf	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

email <i>address</i>	The owner's email address.
name <i>name</i>	The owner's name.
phone <i>number</i>	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 $\verb"owner-info"$ command displays information about the owner of the everRun system.

pm-clear-mtbf

Usage

avcli pm-clear-mtbf

Description

The ${\tt pm-clear-mtbf}$ 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

port <i>number</i>	The port number.
username <i>name</i>	The user's name.
password password	The user's password.
host	The host name.

Examples

\$ avcli --port 8080 proxy.my-domain.com \$ avcli --port 8080 --username user --password secret proxy.mydomain.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.

removable-disk-info

Usage

```
avcli removable-disk-info
```

Description

The removable-disk-info command displays information about USB flash drives that can be mounted into VMs.

In the output, information about each USB flash drive begins with the name (for example,

removabledisk:036). The name is the device ID for commands such as vm-attach-usb-storage.

Examples

The following is the command with sample output:

```
[root@node0 ~]# avcli removable-disk-info
Removable Disks:
removabledisk:o36:
   -> Description: : Imation Nano Pro
   -> Size: : 7739768832 bytes
   -> Vendor: : Imation
   -> Vendor ID: : 0718
   -> Model: : Nano Pro
   -> Product ID: : 063d
   -> Attached to node: : node0
Local node0 information:
        -> Device Path: : /dev/sdc
        -> USB Port: : 3-7.1:1.0
        -> USB Bus Number: : 3
        -> USB Device Number: : 6
```

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. Specifically, the command performs the following actions:

- Enables and disables SNMP requests.
- Enables and disables SNMP traps.
- Specifies the port to use for SNMP traps.
- Specifies the SNMP community.
- Adds SNMPv1 and SNMPv2 recipients.

enable-requests	Enable SNMP requests. If you do not specify the option, requests are disabled.
enable-traps	Enable SNMP traps. If you do not specify this option, traps are disabled. When you enable traps, you must spe- cify one or more recipients.
community <i>name</i>	The name of the SNMP community.
port <i>number</i>	The port to use for SNMP. The default is 162.
recipients <i>recipient</i>	The list of hosts to which to send traps using SNMP version 2c.
recipients-v1 <i>recipient-v1</i>	The list of hosts to which to send traps using SNMP version 1.

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 the configuration of all SNMP versions.

storage-group-create

Usage

avcli storage-group-create storage-group-name...

Description

The storage-group-create command creates one or more storage groups.

storage-group-name	The name of the storage group to be created.
--------------------	--

storage-group-delete

Usage

avcli storage-group-delete storage-group-name...

Description

The storage-group-delete command deletes one or more storage groups.

storage-group-name	The name of the storage group to be deleted.
--------------------	--

storage-group-info

Usage

```
avcli storage-group-info [--disks] [--volumes] [storage-
group...] [--orphan]
```

Description

The storage-group-info command displays information about all storage groups, or optionally, only about those specified.

disks	Show the logical disks belonging to a storage group.
volumes	Show the volumes using a storage group.
storage-group	One or more storage groups about which to display inform- ation.
orphan	Show the logical disks that are not part of any storage group.

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.

disks	Show the logical disks belonging to a storage group.	
volumes	Show the volumes using a storage group.	
storage-group	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

timezone 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.

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

true	Enables VMs to automatically return to a node that has returned to ser- vice
false	Disables VMs from automatically returning to a node that has returned to service.
reset	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*.

cluster-address <i>IP_</i> <i>address</i>	The IP address of the everRun system.
static	Explicitly sets the values, if specified, for prefix,node0-address,node0- gateway,node1-address,node1- gateway, anddns-servers; otherwise, DHCP sets these values (exceptcluster- address).
prefix <i>prefix</i>	The size of the network prefix. Values, in bits, are 8 (Class A), 16 (Class B), and 24 (Class-C).
node0-address <i>IP_</i> <i>address</i>	The IP address of node0.
node0-gateway IP_ address	The IP address of the node0 gateway.
nodel-address <i>IP_</i> <i>address</i>	The IP address of node1.

nodel-gateway <i>IP_</i> <i>address</i>	The IP address of the node1 gateway.
[dns-servers server_ address]	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). 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.

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.

wait	Wait for the command to complete
-w	wait for the command to complete.

vm-attach-usb-storage

Usage

```
avcli vm-attach-usb-storage --name name_or_OID --deviceId
device Id
```

Description

The vm-attach-usb-storage command attaches the specified USB flash drive to a VM on the active node. The USB flash drive must be plugged in to the VM's active node.

Options

name <i>name_or_OID</i>	The name or OID of the VM.
deviceId <i>device_ld</i>	The device ID of the USB flash drive. Output of the removable-disk-info command includes the device ID of a VM.

Examples

\$ avcli vm-attach-usb-storage --name MyVM --deviceId 063d

The following example includes output.

```
$ avcli vm-attach-usb-storage --name buick1 --deviceId
removabledisk:o36
VM: buick1 vmOID vm:o1808 deviceId: removabledisk:o36
Removable Disks:
removabledisk:o36:
removabledisk:o36:
MATCH:
    removabledisk:o36:
    -> Description: : Imation Nano Pro
    -> Size: : 7739768832 bytes
    -> Vendor: : Imation
```

-> Vendor ID: : 0718 -> Model: : Nano Pro -> Product ID: : 063d -> Attached to node: : node0 Local node0 information:

-> Device Path: : /dev/sdc

- -> USB Port: : 3-7.1:1.0
- -> USB Bus Number: : 3
- -> USB Device Number: : 6

vm-ax-disable

Usage

avcli vm-ax-disable --name name --node node

Description

The vm-ax-disable command disables the VM's instance on a selected PM.

Options

_	-name <i>name</i>	The name or ID of a VM.
-	node <i>node</i>	The name or ID of a PM to disable.

Examples

Disable instance on node1 for a VM named MyVM.

\$ avcli vm-ax-disable --name MyVM --node 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 a selected PM.

Options

name <i>name</i>	The name or ID of a VM.
node node	The name or ID of a PM to enable.

Examples

Enable instance on node0 for a VM named MyVM.

\$ avcli vm-ax-enable --name MyVM --node node0

vm-boot-attributes

Usage

```
avcli vm-boot-attributes --priority priority --application-
start-time minutes [--autoStartMode autoStartMode] [vm...]
```

Description

The vm-boot-attributes command sets the boot attributes for the specified VMs.

Options

priority <i>priority</i>	The boot priority; values are 1 to 1000.
application-start-time <i>minutes</i>	The estimated start time of the VM and applic- ation, in minutes. The minimum value is one minute.
autoStartMode <i>autoStartMode</i>	The automatic start-up mode of the VM. Valid values are: • last (default): Use the last value. • on: Turn on automatic start-up. • off: Turn off automatic start-up.
vm	One or more VMs whose boot attributes are being set.

```
$ avcli vm-boot-attributes --priority 1 --application-start-time
1 vm1
$ avcli vm-boot-attributes --priority 1 --application-start-time
1 vm:o100
$ avcli vm-boot-attributes --priority 1 --application-start-time
1 --autoStartMode on vm1
```

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

iso <i>iso</i>	The ISO image to boot.
wait	Wait for the VM to boot.
vm	One or more VMs that are being started.

- \$ 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]
[--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.

source-vm <i>source</i>	The name or ID of the source VM.
name <i>name</i>	The name of the VM to create.
description " <i>description</i> "	The description for the new VM.
cpu <i>number</i>	The number of virtual CPUs to assign to the VM.
memory <i>memory</i>	The amount of memory, in megabytes, to assign to the VM.
availability <i>level</i>	The availability level, high availability (ha) or fault tolerant (ft).
copy-volumes <i>volumes</i>	The list of volumes to copy to the new VM. Specify volumes by configuration name or ID, boot volume first. To copy all of the volumes from the source VM to the new VM with the default values, keep this parameter as blank. A <i>volume</i> is made up of five components, separated by

	commas:
	 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).
	 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.
	The list of volumes to attach to this VM. A <i>volume</i> is made up of five components, separated by commas:
add-volumes <i>volumes</i>	 Size of the volume; required. By default, the volume size is specified in mega- bytes, 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).
keep-volumes <i>volumes</i>	The idle volumes to be attached to the new VM. Specify a volume either by name or ID.
interfaces <i>networks</i>	The list of networks to attach to the VM. Specify a net- work only once. The attached network must not be private.
storage-group <i>group</i>	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 stor- age group is configured with 4 kB sectors, ensure that the guest OS supports the 4 kB sector size.
no-auto-start	If set, the VM is not started after the copy has finished.

Examples

Copy a VM named vm: 02046 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, 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 --no-autostart

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] --
volumes volumes [--wait]
```

Description

The vm-create command creates a new VM.

name <i>name</i>	The name of the VM to create.
cpu <i>number</i>	The number of virtual CPUs to assign to the VM.
memory <i>memory</i>	The amount of memory, in megabytes, to assign to the VM.
cdrom <i>cd-name</i>	The CD-ROM from which you initially boot the VM. You cannot specify this option withkickstart or remote-file-path.
kickstart template	The kickstart template to use when booting the VM. You cannot specify this option withcdrom or remote-file-path.
remote-file-path path	A remote ISO repository to use when booting the VM. You cannot specify this option withcdrom or kickstart.
remote-type <i>type</i>	The type of remote ISO repository specified in the – –

	remote-file-path option . Valid options are samba or nfs.
remote-username <i>username</i>	The user account to specify for access to the remote ISO repository specified in theremote-file- path option. Required for samba repositories.
remote-password password	The user password to specify for access to the remote ISO repository specified in theremote-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> , <i>MAC address</i>	The list of networks to attach to the VM. Specify a net- work only once. The attached network must not be private. Optionally, specify the MAC address after the network name.
disabled-interfaces <i>net-</i> <i>works, MAC address</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. Optionally, spe- cify the MAC address after the network name.
storage-group group	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.
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: • Size of the volume; required. By default, the volume size is specified in mega- bytes, 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).
	 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.
wait -w	Wait for the command to complete.

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, 1024 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 1024 \
--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 \
--boot-type bios --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-7.7-x86_64-minimal.iso \
--availability ha --interfaces network0 \
--volumes 20GB,Pool-0001,vm001_volu_boot,qcow2 1024,Pool-
0002,\
vm001 volu data,qcow2,4096
```

Create the HA VM vm001 with a CPU, 2048 MB of memory, a 1,024 MB volume, and attached to network0 with MAC address 00:04:fc:40:60:55.

```
$ avcli vm-create --name vm001 --cpu 1 --memory 2048 \
--cdrom linux.iso --availability ha \
--interfaces network0,00:04:fc:40:60:55 --volumes 1024
```

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 ${\tt vm-create-from-snapshot}$ command creates a new VM from a VM snapshot.

vm-snapshot-oid oid	The vm snapshot OID from which to create the VM.
name <i>name</i>	The name of the VM to create.
cpu <i>number</i>	The number of virtual CPUs to assign to the VM.
memory <i>memory</i>	The amount of memory, in megabytes, to assign to the VM.
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.
storage-group group	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.
volumes <i>volumes</i>	Restrict the included volumes to those specified; oth- erwise, all volumes will be created. Specify volumes by

	configuration name or ID, boot volume first.
volume-prefix <i>prefix</i>	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 ocean_boot, and you specifyvolume-prefix new, the imported volume would be new-ocean_ boot.
no-auto-start	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 centosboot centos-data --volume-prefix minimal --no-auto-start

vm-delete

Usage

avcli vm-delete [--volumes] [--wait] vm...

Description

The vm-delete command deletes the specified VMs, the snapshots belonging to those VMs, and optionally, the volumes attached to the VMs.

Options

volumes	Delete the volumes attached to the VM.
wait -w	Wait for the command to complete.
vm	One or more VMs to be deleted.

Examples

avcli vm-delete vm1 avcli vm-delete --volumes vm1 avcli vm-delete --volumes vm1 vm2

vm-device-config-info

Usage

```
avcli vm-device-config-info
```

Description

The vm-device-config-info command displays configuration information about VM devices.

For Disable insert of media to all VMs, the command displays false (the default) when insertion of media (for example, virtual CDs) is enabled and true when insertion of media is disabled.

```
For Disable attach of USB devices to all VMs, the command displays false (the default) when attaching USB devices (for example, a flash drive) is enabled and true when attaching USB devices is disabled.
```

- \$ avcli vm-device-config-info
 - VM device configuration:
 -> Disable insert of media to all VMs : false
 - -> Disable attach of USB devices to all VMs : false

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 .

Options

Π

path pathname	A pathname relative to the export mount point to where the expor- ted OVF is written.	
volumes <i>name</i>	Restrict the exported volumes to those specified; otherwise, all volumes will be created. Specify volumes by configuration name or ID, boot volume first.	
wait	Wait for the export operation to complete. Specify this option to view the export progress.	
force	Force the VM to be exported, even if it is still running.	
vm-name	Specify the name of the VM to be exported.	

```
$ 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.ovf [--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] [--image-format format]
```

Description

The vm-import command imports a VM from an OVF VM archive file.

archive <i>filename.ovf</i>	The OVF file archive to import.
no-auto-start	Do not start the VM after the import has finished.
cpu <i>number</i>	The number of CPUs to assign to the VM. This defaults to the value in the archive.
memory <i>size</i>	The size of memory, in megabytes, to assign to the VM. This defaults to the value in the archive.
name <i>vm-name</i>	The name to assign to the VM. This defaults to the value in the archive.
storage-groups groups	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.
interfaces <i>networks</i>	The list of shared networks to assign to the VM's inter- faces. By default, values in the archive or available shared networks are assigned.

remap-volumes	First attempt to remap all volumes to the shared-mirrors as defined in the archive, after that thevolumes andstorage-groups rules are in effect.
volumes <i>volumes</i>	Import only these volumes. By default, all available volumes from the OVF are imported.
volume-prefix prefix	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 ocean_boot, and you specifyvolume- prefix new, the imported volume would be new- ocean_boot.
data	Import data only for the specified volumes.
force	When the OVF file is missing the isBootable flag (a known issue for Windows XP), assume that the VHD pointed to by OVF is the bootable one.
silent	Suppress output.
dry-run	Show the interface to the shared network and volume-to- storage-group assignments without actually importing or restoring a VM.
throttle <i>amount</i>	 Slow down the import/export operation. Valid values are: none: Do not use throttling. This is the default value. low: Slow down by about 25%. medium: Slow down by about 50%. high: Slow down by about 75%.
use-https	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.
protection-level <i>level</i>	The protection level to assign to the VM. Valid options are HA and FT (default).
image-format format	The image format for all disk volumes of the VM. Valid values are gcow2 and raw (default).

```
$ avcli vm-import --archive vml.ovf
$ avcli vm-import --name myVM --throttle low --archive vml.ovf
$ avcli vm-import --cpu 2 --memory 1024 --archive vml.ovf
$ avcli vm-import --interfaces network0 network1 --archive
vml.ovf
$ avcli vm-import --remap-volumes --archive vml.ovf
$ avcli vm-import --remap-volumes sm-0000 sm-0001 --archive
vml.ovf
$ avcli vm-import --volumes boot_vol vol3 --data vol3 --archive
vml.ovf
$ avcli vm-import --name myVM --protection-level HA --archive
vml.ovf
$ avcli vm-import --name myVM --protection-level HA --archive
vml.ovf
```

vm-info

Usage

```
avcli vm-info [vm...]
```

Description

The vm-info command displays information about all VMs, or optionally, about specific VMs.

Options

vm	One or more VMs for which to display information.
----	---

- \$ avcli vm-info
- \$ avcli vm-info vm1
- \$ avcli vm-info vm1 vm:o100

vm-media-insert-disable

Usage

avcli vm-media-insert-disable

Description

The vm-media-insert-disable command disables the function of inserting media (for example, a virtual CD) into all VMs. (To disable the function of inserting a USB device, use vm-usb-attach-disable.)

Only users whose role is Administrator can issue this command.

vm-media-insert-enable

Usage

avcli vm-media-insert-enable

Description

The vm-media-insert-enable command enables the function of inserting media (for example, a virtual CD) into all VMs. (To enable the function of inserting a USB device, use vm-usb-attach-enable.)

Only users whose role is Administrator can issue this command.

vm-network-disable

Usage

```
avcli vm-network-disable --name name --node node --networks
```

Description

The vm-network-disable command disables a VM's networks on a selected node.

Options

name <i>name</i>	The name or ID of a VM.
node <i>node</i>	The name or ID of a PM on which to disable the networks.
networks <i>networks</i>	The names or IDs of networks to disable.

Examples

Disable net2 on node1 for a VM named MyVM.

```
\ avcli vm-network-disable --name MyVM --node nodel \setminus
```

```
--networks net2
```

vm-network-enable

Usage

```
avcli vm-network-enable --name name --node node --networks networks
```

Description

The vm-network-enable command enables a VM's networks on a selected node.

Options

name <i>name</i>	The name or ID of a VM.
node <i>node</i>	The name or ID of a PM on which to enable the networks.
networks <i>networks</i>	The names or IDs of networks to enable.

Examples

Enable net2 on node1 for a VM named MyVM.

```
\ avcli vm-network-enable --name MyVM --node nodel \setminus
```

--networks net2

vm-poweroff

Usage

avcli vm-poweroff [vm...] [--wait]

Description

The ${\tt vm-poweroff}$ command powers off the specified VMs.

Options

vm	One or more VMs to power off. Specify a VM either by name or ID.
wait -w	Wait for the command to complete.

- \$ 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 ${\tt vm-poweron}$ command powers on the specified VMs.

Options

vm	One or more VMs to power on. Specify a VM either by name or ID.
wait -w	Wait for the command to complete.

- \$ 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] [--detach-boot-volume] [--attach-boot-volume name]
```

Description

The vm-reprovision command reprovisions the specified VM.

name <i>name</i>	Specify the VM to reprovision. Reprovision only one VM at a time. Specify a VM either by name or ID.
new-name <i>name</i>	Specify a new name for the VM.
description "description"	Specify a description for the VM.
cpu <i>number</i>	The number of virtual CPUs. This defaults to the VM's current amount.
memory <i>size</i>	The size of memory, in megabytes. This defaults to the VM's current amount.
addVolumes <i>volumes</i>	 The list of volumes to attach to this VM. A <i>volume</i> is made up of five components, separated by commas: 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). 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.
	B as the sector size.
deleteVolumes <i>volumes</i>	I he list of volume(s) that are currently attached to the spe- cified VM to be deleted. Specify a volume either by name or ID.
keepVolumes <i>volumes</i>	The list of volume(s) that are currently attached to the spe- cified 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.
interfaces <i>networks,</i> MAC address	The list of networks to attach to the VM. Specify a network only once. The attached network must not be private. Optionally, specify the MAC address after the network name.
disabled-interfaces <i>networks, MAC address</i>	A list of networks to attach to the VM, but which should not be enabled. Specify a network only once. The attached net- work must not be private. Optionally, specify the MAC address after the network name.

detach-boot-volume	Detaches the VM's boot volume.
	Specify the name of a new boot volume for this VM. If the
attach-boot-	VM already has an attached boot volume, you must also
volume <i>name</i>	<pre>specifydetach-boot-volume; otherwise, the</pre>
	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 $vm001_vo10$.

```
$ avcli vm-reprovision --cpu 1 --memory 512 --interfaces
network0 \
```

--addVolumes 1024, Pool-0001, vm001 vol0 --name vm1

Reprovision VM vml, and then delete the volumes volume: 0411, data-vml, 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-bootvolume 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.ovf [--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 ${\tt vm-restore}$ command restores a VM from an OVF file.

archive <i>filename.</i> ovf	The OVF file archive to restore.
no-auto-start	Do not start the VM after the restore has finished.
cpu <i>number</i>	The number of CPUs to assign to the VM. This defaults to the value in the archive.
memory <i>size</i>	The size of memory, in megabytes, to assign to the VM. This defaults to the value in the archive.
name <i>vm-name</i>	The name to assign to the VM. This defaults to the value in the archive.
storage-groups <i>groups</i>	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.
interfaces <i>networks</i>	The list of shared networks to assign to the VM's inter- faces. By default, values in the archive or available shared networks are assigned.
volume-prefix prefix	Adds the specified <i>prefix</i> to the beginning of the newly

	<pre>imported volume names to prevent conflicts with existing volumes on the system. For example, if a source volume is ocean_boot, and you specify - -volume-prefix new, the imported volume would be new-ocean_boot.</pre>
data	Restore data only for the specified volumes.
silent	Suppress output.
dry-run	Show the interface to the shared network and volume- to-storage-group assignments without actually restor- ing a VM.
throttle	 Slow down the operation. Valid values are: none: Do not use throttling. This is the default value. low: Slow down by about 25%. medium: Slow down by about 50%. high: Slow down by about 75%.
use-https	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.

- \$ avcli vm-restore --archive vm1.ovf
- \$ 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 ${\tt vm-shutdown}$ command shuts down the specified VMs.

Options

vm	One or more VMs to shut down. Specify a VM either by name or ID.
wait -w	Wait for the command to complete.

- \$ 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
"description"] [--desire] [--require] vm-name
```

Description

 $The {\tt vm-snapshot-create} \ command \ creates \ a \ VM \ snapshot.$

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.

volumes no-data	The names of volumes to include in the snapshot. By default, all volumes are included in the snapshot unless you specify volumes with individual volume names or you specifyno-data. If you specifyno-data, no volumes are included in the snapshot. These arguments are mutually exclusive.
description " <i>description</i> "	The user-specified description for this snapshot.
desire	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 require). Values are crash and application (the default value).
require	The minimum consistency level needed to declare the snapshot a success. Values are crash and application (the default value).
---------	---
vm-name	The VM's ID.

Examples

\$ avcli vm-snapshot-create --volumes volume:o100 volume:o101
--description "This is the snapshot description" --name
snapshot_name 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-createdisable

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-createenable

vm-snapshot-delete

Usage

avcli vm-snapshot-delete snapshot...

Description

The ${\tt vm-snapshot-delete}$ command deletes the specified snapshots.

Options

snapshot 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 snapshot of 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.

Options

П

wait	Wait for the export operation to complete. Specify this option to view the export progress.
silent	Suppress progress output.
volumes <i>volumes</i>	Restrict the exported volumes to those specified; otherwise, all volumes will be created. Specify volumes by configuration name or ID, boot volume first.
path pathname	A pathname relative to the export mount point to where the expor- ted OVF is written.
snapshot	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

snapshot 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



Examples

```
$ avcli vm-unlock vm1
$ avcli vm-unlock vm:o100
```

vm-usb-attach-disable

Usage

avcli vm-usb-attach-disable

Description

The vm-usb-attach-disable command disables the function of attaching USB storage devices to all VMs. (To disable the function of inserting a virtual CD, use vm-media-insert-disable.) Only users whose role is Administrator can issue this command.

vm-usb-attach-enable

Usage

avcli vm-usb-attach-enable

Description

The vm-usb-attach-enable command enables the function of attaching USB storage devices to all VMs. (To enable the function of inserting a virtual CD, use vm-media-insert-enable.)

Only users whose role is **Administrator** can issue this command.

vm-volume-disable

Usage

```
avcli vm-volume-disable --name name --node node --volumes volumes
```

Description

The vm-volume-disable command disables a VM's volumes on a selected node.

Options

name <i>name</i>	The name or ID of a VM.
node <i>node</i>	The name or ID of a PM on which to disable the volumes.
volumes <i>volumes</i>	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 --node node1 \
--volumes abba1-data volume:o2249
```

vm-volume-enable

Usage

avcli vm-volume-enable --name name --node node --volumes volumes

Description

The vm-volume-enable command enables a VM's volumes on a selected node.

Options

name <i>name</i>	The name or ID of a VM.
node <i>node</i>	The name or ID of a PM on which to enable the volumes.
volumes <i>volumes</i>	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 --node nodel \setminus
```

--volumes volume:02249

volume-info

Usage

```
avcli volume-info [volume...]
```

Description

The volume-info command displays information about all volumes, or optionally, only about those specified.

Options

volume	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

new-size <i>size</i>	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).
volume	The volume to be resized.

Examples

avcli volume-resize --new-size 79G boot-airplane1

11

Chapter 11: System Reference Information

See the following topics for reference information

- Compatible Guest Operating Systems
- Physical Machine System Requirements
- Important Physical Machine and Virtual Machine Considerations
- Accessing Knowledge Base Articles
- Creating a SplitSite Configuration

Compatible Guest Operating Systems

The following are compatible as guest operating systems for virtual machines (VMs) running on everRun systems.

Operating System	Version	Boot Firmware Interface
CentOS 7	CentOS 7.0 ¹ , 7.1, 7.2, 7.3, 7.4, 7.5 (all 64-bit)	BIOS

¹CentOS 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.

Operating System	Version	Boot Firmware Interface
CentOS 6	CentOS 6.7, 6.8, 6.9 (all 64-bit)	BIOS
Microsoft Windows Server 2016 (Essentials, Standard, Datacenter)	64-bit	BIOS
Microsoft Windows Server 2012 (Foundation, Essen- tials, Standard, Datacenter)	64-bit, 64-bit R2	BIOS
Microsoft Windows Small Business Server 2011 (Standard, Essential, Premium Add-On)	64-bit	BIOS
Microsoft Windows Server 2008 (Web, Small Busi- ness, Standard, Enterprise, Datacenter)	32-bit, 64-bit R2 only	BIOS
Microsoft Windows Server 2003 (Enterprise)	32-bit R2 SP2 ¹	BIOS
Microsoft Windows 10 Desktop	64-bit	BIOS
Microsoft Windows 8.1 Desktop (Enterprise)	64-bit	BIOS
Microsoft Windows 8 Desktop (Enterprise)	64-bit	BIOS
Microsoft Windows 7 Desktop	32-bit, 64-bit	BIOS
Oracle Linux 7.3 Red Hat Compatible (RHC) Kernel	x86-64	BIOS
Oracle Linux 6.8 RHC Kernel	x86-64	BIOS

¹For specific installation and migration procedures, see Creating a New Windows Server 2003 Virtual Machine and Migrating a Windows Server 2003 VM to an everRun 7.x System.

Operating System	Version	Boot Firmware Interface
Red Hat Enterprise Linux 7 (Workstation, Server)	Red Hat 7.0 ¹ , 7.1, 7.2, 7.3, 7.4, 7.5 (all 64-bit)	BIOS
Red Hat Enterprise Linux 6 (Workstation, Server)	Red Hat 6.7, 6.8, 6.9 (all 64-bit)	BIOS
SUSE Linux Enterprise Server	SLES 11 SP3 64-bit SLES 11 SP4 64-bit SLES 12 SP2 64-bit SLES 12 SP3 64-bit	BIOS
Ubuntu	12.04 LTS, 14.04 LTS, 16.04 LTS (all 64-bit) 17.10.1 Server 64-bit 18.04 Server 64-bit	BIOS

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	Notes
Processors/CPUs: Intel [®] Xeon [®] Gold 61 <i>XX</i> , 51 <i>XX</i> (Skylake) Intel Xeon Silver 41 <i>XX</i> (Skylake)	1	2	

¹RHEL 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.

Physical Device	Minimum	Maximum Tested	Notes
Intel Xeon Bronze 31 <i>XX</i> (Skylake)			
Intel Xeon E5-2 <i>XXX</i> v4 (Broadwell)			
Intel Xeon E5-2 <i>XXX</i> v3 (Haswell)			
Intel Xeon E5-2 <i>XXX</i> v2 (Ivy Bridge)			
Intel Xeon E5- <i>2XXX</i> (Sandy Bridge)			
Intel Xeon E5-1 <i>XXX</i> v4 (Broadwell)			
Intel Xeon E5-1 <i>XXX</i> v3 (Haswell)			
Intel Xeon E5-1 <i>XXX</i> v2 (Ivy Bridge)			
Intel Xeon E5-1 <i>XXX</i> (Sandy Bridge)			
Intel Xeon E3-1 <i>XXX</i> v6 (Kaby Lake)			
Intel Xeon E3-1 <i>XXX</i> v5 (Skylake)			
Intel Xeon E3-1 <i>XXX</i> v4 (Broadwell)			
Intel Xeon E3-1 <i>XXX</i> v3 (Haswell)			
Intel Xeon E3-1 <i>XXX</i> v2 (Ivy Bridge)			

Physical Device	Minimum	Maximum Tested	Notes
Intel Xeon E3-1 <i>XXX</i> (Sandy Bridge)			
Number CPU Sockets per Physical Machine	1	2	
Physical Memory	8 GB	384 GB	
Internal Disk Count per Phys- ical Machine	2	24	Minimum 2 drives per PM for FT. VM disks/volumes are rep- licated on both PMs.
Total Disk Capacity	36 GB	9.4 TB	
Management ENET Ports	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 Man- agement 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
- Virtual Machine Recommendations and Limits
- Combined Virtual Machine Maximums
- Important Considerations

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 <i>n</i> vCPUs	n+2 (typical)
Each HA guest with <i>n</i> 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

٥

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	8
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	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
Recommended minimum configuration for physical mach • One logical volume protected by RAID1, RAID 5, RAID 10 or • Two non-RAID or RAID 0 volumes. When using multiple volumes per RAID set, the RAID set a type that provides redundancy, such as RAID1, RAID5 RAID10.	Recommended minimum configuration for physical machines:
	 One logical volume protected by RAID1, RAID 5, RAID 6, or RAID 10
	or
	Two non-RAID or RAID 0 volumes.
	When using multiple volumes per RAID set, the RAID set should set
	a type that provides redundancy, such as RAID1, RAID5, or
	RAID10.
QCOW3 (QCOW2v3)	QCOW2 refers to both QCOW2 and QCOW3 files in descriptions of
	everRun systems. By default, an everRun system creates QCOW3
	mes (-i qcow2 -o compat=1.1).
USB CD/DVD Drive	USB CD/DVD drives are supported on all platforms for everRun installation.
Direct-Attach Tape Drives	Access to direct-attach tape drives by the guests is not
	supported. Stratus recommends using network-attach tape drives.
Console Connectivity	Each PM's text console is available for a CentOS operating system. However, VGA mode is not supported; that is, the PM must be at run-

Feature	Comment
	level 3 and cannot run at run-level 5. See "System Management" below.
SSD Support	everRun supports solid state drives according to the storage con- troller vendor specifications.
System Management	everRun system management does not work at run-level 5.
Volumes	For exporting, importing, or restoring a volume, the maximum volume size is 2TB.

Creating a SplitSite Configuration

П

This topic and its subtopics describe how to create a SplitSite configuration. For general information about quorum servers, see Quorum Servers as well as SplitSite and Quorum Service.

Note: Before you create a SplitSite configuration, read this topic and all of its subtopics and then plan your SplitSite configuration, as described in the topics. Create the configuration only after you are certain that your planned configuration complies with the information in this topic and its subtopics.

A SplitSite configuration exists if either of the following is true:

- The two nodes of the system are connected using network infrastructure rather than direct cables.
- The length of the A-Link (direct connect) cables connecting the two nodes is greater than 10m (for example, in two separate buildings within a campus).

These configurations provide better disaster tolerance and hardware redundancy as well as redundancy of physical computer rooms and the buildings containing them.

Stratus recommends that a SplitSite configuration include a third computer, which is a quorum server. The quorum server is located in a physical location that is removed from the physical location of both node0 and node1.

ñ

ñ

Note: This topic and its subtopics describe a SplitSite configuration with a quorum server. Stratus highly recommends that a SplitSite configuration include a quorum server. If you want to consider creating a SplitSite configuration without a quorum server, access the Knowledge Base to search for the article *Considerations if deploying SplitSite without quorum* (KB-9683), and also contact your authorized Stratus service representative. For information about accessing Knowledge Base articles, see Accessing Knowledge Base Articles.

Because of the geographic separation of these physical machines, creating a SplitSite configuration requires careful planning of component placement and more complex networking topologies.

The topics below describe how to create a SplitSite configuration. To perform the procedures in the topics, you should be familiar with everRun software and the hardware it runs on, and you should be familiar with the network infrastructure of your system and its location.

Note: These topics cannot describe every vendor and model of network switches, routers, and other hardware. Consult the documentation that pertains to your infrastructure if you need more information about how to configure it according to the requirements in these Help topics.

- Creating the Configuration
- Meeting Network Requirements
- Locating and Creating the Quorum Server
- Completing the Configuration
- Understanding Quorum's Effect on System Behavior

The following table lists and defines terms associated with creating a SplitSite configuration.

Term	Meaning
Active node	The node where a guest VM is currently running. Each guest VM may have a different active node. The opposite of <i>active</i> is standby (see Standby node).
A-Link	Availability link. A direct network connection between the two computers that form an everRun system. (The computers of a system are also referred to as <i>physical machines</i> (PMs) or <i>nodes</i> .) A-Links must be con-

	nected point-to-point, and A-Link traffic cannot be routed. A everRun sys- tem requires two A-Links. On some systems, these connections have blue and yellow cables (and ports). You can use VLAN connections for A-Links in a distributed local site installation (see VLAN).
Alternate quorum server	The alternate quorum server is used when the preferred quorum server is not available (see Preferred quorum server).
AX	The container layer that resides within the everRun system and controls the behavior of the guest VM. AX is responsible for keeping a VM syn- chronized between the active node and the standby node. Each VM has its own AX pair (see VM, Active node, and Standby node)
Business network (ibiz)	A network connection from the everRun system to a LAN that also has other traffic that can include management messages as well as traffic for applications and other clients and servers. The everRun system typically has two ports for business network connections. Business networks can be assigned to one or more guest VMs for their use, or to no guest VMs. You must connect the first business network (ibiz0) to a LAN so that you can manage the system from a web browser.
Fault	Any potential degradation in a system's ability to execute a guest VM (see VM). Disk failure, network loss, or power outage are all examples of faults detected by the system.
Node0 and node1	The two computers that form the everRun system are labeled internally as node0 and node1. (These computers are also sometimes referred to as physical machines or PMs.) The choice of node0 and node1 is arbitrary and is made when the system is configured for the first time. Constant traffic flowing between node0 and node1 communicates state information for the system as well as for each running guest VM (see VM).
Preferred quorum server	The preferred quorum server is used when it (the preferred quorum server) is available. If the preferred quorum server is not available, the alternate quorum server (if it exists) is used (see Alternate quorum server).

Primary node	When the system's computers are paired, only one computer responds to management messages. This computer is the primary node. The System IP address, which is assigned when the system is initially installed, applies to the primary node. The primary node can switch between node0 and node1 as different fault conditions occur (see Fault). Note that the primary node is not necessarily the active node for a guest VM (see Active node and VM).
priv0	A network for private management traffic between the two nodes. For more information, see A-Link and Private Networks.
Quorum server	A third computer that helps arbitrate which AX should be active for each guest VM (see Active node and VM). Correct use of a quorum server is the only guaranteed way to prevent split-brain conditions (see Split-brain).
RTT	Round-trip time. The elapsed time required for a network message to travel from a starting point to a destination and back again. The time is typically measured in milliseconds (ms).
Split-brain	The condition that occurs when both AX's of a guest VM's AX pair are act- ive simultaneously, which produces divergent copies of data within each active guest (see AX and VM). Split-brain can occur when all com- munication paths between node0 and node1 are disconnected (see Node0 and node1). Using the quorum service prevents a split-brain condition (see Quorum server).
SplitSite	 A SplitSite configuration exists if either of the following is true: The two nodes of the everRun system are connected using network infrastructure rather than direct cables. The length of the A-Link (direct connect) cables connecting the two nodes is greater than 10m (for example, in two separate buildings within a campus). A SplitSite configuration is typically used to provide better disaster tol-

	erance, at the expense of more network setup and more extensive con- figuration options. A SplitSite configuration requires a third computer, which is a quorum server (see Quorum server).
Standby node	The node that is not the active node for a guest VM. The standby node is kept synchronized through AX communications via A-Link connections (see AX and A-Link). The AX pair for each guest VM determines which node is active and which is standby (see Active node).
System management	System management is the layer within everRun software that is respons- ible for maintaining the overall state of the system. Determining which node is primary is part of system management (see Primary node). Sys- tem management is also responsible for displaying information within the everRun Availability Console.
UPS	Uninterruptable power supply. An external battery backup for electrical equipment that prevents short power outages from affecting availability.
VLAN	Virtual LAN. A VLAN is a set of devices on one or more LANs that are con- figured to communicate as if they were attached to the same cabled net- work, when in fact they are located on different LAN segments. VLANs are configured at the network infrastructure level, not within the everRun sys- tem. In aSplitSite configuration, the A-Link connections are implemented as isolated VLANs (see A-Link).
VM	Virtual Machine (also referred to as a guest). A system typically has one or more VMs (or guests) allocated and running applications via guest oper- ating systems.

Creating the Configuration

To create a SplitSite configuration, first consider the configuration of a typical everRun system configuration and the VLAN requirements of a SplitSite configuration. Then, observe a well-planned SplitSite configuration, which includes a quorum server, and become familiar with the configuration's VLAN requirements. You must also become familiar with the entire process of installing a typical everRun system and then creating a SplitSite configuration. The sections below provide this information.

A Typical everRun System

In a typical everRun system configuration, two PMs are directly connected by at least two network cables for A-Links. One A-Link typically serves as the private network (priv0). The two PMs have additional network connections for business networks, which the everRun Availability Console and guest VMs hosted by the system use. The following figure illustrates a typical configuration. (In the figures, the port labels A2, A1, P2, and P1, are an example, though most everRun systems do not have port labels.)



The physical distance between the PMs in a typical configuration is limited by the length of a single A-Link network cable, which is approximately 33 ft (10m). This distance may be significantly shorter when the physical environment and ambient electrical noise is accounted for.

A SplitSite Configuration With a Quorum Server

A well-planned SplitSite configuration consists of the two nodes at two different locations, and a third computer that runs the quorum service at a third location. All of these computers are networked together with appropriate network switching equipment, so that no single point of failure exists within the SplitSite configuration. The following figure illustrates such a configuration, which includes node0 at Site A, node1 at Site B, and the quorum server at Site C.





SplitSite VLAN Requirements

The A-Link connections between switch A and switch B require a VLAN configuration on the switches. A-Link traffic is not routable, and the connection should emulate a single long network cable. Each A-Link must be isolated on its own VLAN.

If you cannot create VLANs between the switching equipment, you can use Ethernet-to-fiber media converters to create a longer fiber connection between the two PMs. However, you should not route the two A-Link fiber connections through the same physical conduit, as this creates a single point of failure.

Additionally, the quorum service computer should not share a switch with either node0 or node1 because a shared switch creates a single point of failure.

See Meeting Network Requirements for more information about the latency requirements of the A-Links and quorum connections.

From Initial Installation to Completing the SplitSite Configuration

When creating an SplitSite configuration, you should first install a typical everRun system, initially without the SplitSite configuration. For simplicity, install the nodes side-by-side. See Getting Started.

After the typical system is operating normally, create the SplitSite the configuration.

- 1. Read Creating a SplitSite Configuration and all of its subtopics, if you have not already done so.
- 2. Install the quorum computer and enable the quorum server. Comply with all information in:
 - A SplitSite Configuration With a Quorum Server
 - SplitSite VLAN Requirements
 - Meeting Network Requirements
 - Completing the Configuration
- 3. Verify that the quorum server has access to both nodes.
- 4. Properly shutdown one node. See Shutting Down a Physical Machine.
- 5. Relocate the shutdown node to the far site.
- 6. Connect the infrastructure. The SplitSite-configuration illustration above shows the connections, which include:
 - The priv0 connection
 - The ibiz0 connection
 - The second A-Link connection
- 7. Power on and (re-)join the nodes.
- 8. Verify the configuration. Ensure that:
 - The shared networks pair properly—In the everRun Availability Console, navigate to the Networks page and ensure that the state of each network is green-checked. If necessary, troubleshoot any infrastructure problems.
 - Quorum connections are remade—In the console, navigate to the Quorum Servers page by clicking Preferences and then Quorum Servers. Ensure that the state of the quorum server is green-checked. If necessary, troubleshoot any infrastructure problems.

- The primary node can shift from node0 to node1, and the console can connect in both configurations—Place each node in Maintenance Mode (see Maintenance Mode).
- (Re-)join the VMs–Migrate the VMs from node to node (see Migrating a Physical Machine or Virtual Machine to a System). Verify the correct network failover of VM networking.
- 10. Assess the status of network and validate Ethernet failover (see The Networks Page).

Meeting Network Requirements

This topic describes the network requirements and considerations of A-Links, business networks, private networks, the quorum server connections, and the management network for a successful SplitSite configuration. (For general information about these networks, see Network Architecture.)

Prerequisite: Plan and create a SplitSite configuration by first reading Creating a SplitSite Configuration and following its instructions, if you have not already done so.

A-Link network connections must meet the following requirements:

- The A-Links use IPv6 addressing.
- A-Links can be dedicated point-to-point fiber connections. If they are not, they must be configured on a VLAN, where each A-Link is connected on its own VLAN.
- FT VMs require less than 2ms RTT A-Link latency.
- HA VMs require less than 10ms RTT A-Link latency.
- Each A-Link must have enough bandwidth to meet the needs of all VMs on the system. Whenever possible, provide, per A-Link, at least one NIC of at least 1Gb and full duplex; use 10 Gb if possible.
- When planning your network infrastructure, you need to account for the uplink bandwidth between the switch and the network backbone across all the ports in use on that switch.
- Do not use a common card (multiport NIC) for both A-Links.
- If these requirements are not met, guest VMs may run more slowly due to limited synchronization bandwidth between the two nodes.

The first business network (ibiz0) is used for communication between the nodes and to the quorum server. The ibiz0 network must meet the following requirements:

- The two nodes must be on the same subnet.
- The network must allow IPv6 multicast traffic between the two nodes.

П

• The two nodes can access the quorum server using IPv4 network addressing.

Private network connections (priv0 through privn) must meet the following requirements:

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

Network connections for the quorum server must meet the following requirements:

- · Access to the quorum service must be provided using ibiz0, using IPv4 network addressing.
- Two UDP ports must be open and available for communication between the nodes and the quorum service, including in the firewalls. By default, these ports are 4557 and 4558. If you want to change these ports, see Configuring the Quorum Service Port (on the quorum computer) and Configuring the Quorum Server Within the everRun Availability Console.
- Latency between an everRun node and the quorum computer should be less than 500ms RTT.
- Throughput is not an important consideration. 10Mb Ethernet, or even T1 bandwidth is adequate.
- Quorum computers are common to all VMs on the same everRun system.
- Quorum computers may be shared among many everRun systems.
- Quorum computers must never be implemented as a VM on the same everRun system that uses it.
- Use different network infrastructure, don't share. An everRun node should not depend on a gateway or switch/router on the partner node site for sustained access to a quorum services computer.

Note: Do not implement the quorum service as a guest VM on a different pair of nodes; a failure on those nodes would cause the VM running the quorum service to failover, which would create unnecessary complications for network topology and fault management. Additionally, a second quorum computer is needed to manage quorum for the everRun system that is running the quorum service.

Management network connections must meet the following requirements:

- By default, the management network is shared with a business network. In this case, all requirements for a business network apply.
- Configure gateways to a business LAN for remote management.

Locating and Creating the Quorum Server

In a well-planned SplitSite configuration, a third computer hosts the quorum service. The quorum service processing requirement is small, so any other existing computer or VM that meets all network and operating requirements can host the quorum service. An effective quorum server depends upon the location of the quorum computer within your network.Stratus recommends configuring two quorum servers, if possible. With two quorum servers, one is the preferred quroum server and the other is the alternate quorum server.

After you have determined an effective location for the quorum computer (and an alternate quorum computer, if desired) and ensured that the computer meets the requirements of the quorum service, you can create the quorum server.

Prerequisite: Plan and create a SplitSite configuration by first reading Creating a SplitSite Configuration and following its instructions, if you have not already done so.

Locating the Quorum Computer

Locate the first quorum computer in a third site within your network, as A SplitSite Configuration With a Quorum Server illustrates. If a third site is not available, locate the quorum computer in a physical location that is different from the physical location of node0 and node1. Locating the quorum computer in a unique site maximizes the chance of the system surviving a problem that causes the loss of both nodes and the quorum computer (for example, a transient electrical, plumbing, or other problem that causes loss of network connectivity).

You should connect the quorum computer to an electrical circuit that is different from the electrical circuit that powers node0 or node1. In addition, you should connect the quorum computer to a UPS unit.

.

Caution: If both AX's lose connectivity with the quorum server, they will attempt to select an alternate quorum server. If no quorum server can be selected, the VM is downgraded to simplex mode, to prevent a split-brain condition if another failure occurs.

If one node shuts down and the VM (AX) on the remaining node cannot reach either the quorum server or its peer, it will shut itself down to avoid the risk of a split-brain condition.

When locating the quorum computer:

- Ensure that the quorum computer does not share a switch (or router) with either node0 and node1.
- Do not use a guest VM within the everRun system to run the quorum service.

See Understanding Quorum's Effect on System Behavior for a description of system behavior and failure modes.

Adding an Alternate Quorum Computer

You can add another quorum computer (with its switch) to your system to create an alternate quorum service. The most common use of an alternate quorum server is when, for example, operating system updates are being applied to the preferred quorum computer. When the preferred quorum computer restarts, the alternate quorum computer is selected and no downgrade occurs. When the preferred quorum is recovered, the selection moves back to the original preferred quorum computer.

When creating a second quorum service, you must follow all of the requirements for the network and quorum placement. If both nodes can communicate with each other and with the same quorum server (either the preferred or alternate quorum server), the system can maintain VM redundancy, even if one quorum connection is lost. Preferred quorum server selection occurs when both nodes have access to each other and to the preferred quorum server. Thus, if the preferred quorum service is lost at the same time a node is lost, the remaining node shuts down the VM even if a second, non-preferred quorum service is available. However, if the preferred quorum server, the selection is noved to the alternate quorum server. Fault handling occurs in a context of the selected quorum server only.

If you create an alternate quorum service, you need to add a second quorum IP address when adding the quorum service in the everRun Availability Console.

Quorum Computer Requirements

You can install quorum service software on any general-purpose computer, laptop, or VM that is running the Windows operating system and that meets these requirements:

- The computer can continually remain powered on and connected to the network such that the ibiz0 network of the everRun system can always access the quorum server.
- The computer has a static IPv4 network address. Do not use DHCP.
- The operating system is Windows Server 2016, Windows Server 2012, Windows Server 2008, Windows 7, or Windows 10; Embedded versions of the Windows OS are not supported.
- A minimum of 100 MB disk space is available.
- Two UDP ports must be open and available for communication between the nodes and the quorum service, including in the firewalls. By default, these ports are 4557 and 4558. To change these ports, see Configuring the Quorum Service Port (on the quorum computer) and Configuring the Quorum Server Within the everRun Availability Console.

Downloading and Installing the Quorum Service Software

After you have determined an appropriate location for the quorum computer, download and install the required software to create the quorum server.

To download and install the quorum server software

- 1. Open the **Downloads** page at https://www.stratus.com/services-support/downloads/?tab=everrun.
- 2. Scroll down to the **Drivers and Tools** section and then click **Quorum Service** to download the quorum server software installer file to the quorum server.
- 3. On the quorum server, double click the installer file.
- 4. Move the downloaded file to an accessible location.
- 5. Log in to the quorum computer.
- 6. Navigate to the quorum service installer and double-click it.
- 7. Follow the prompts to complete the installation.



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

ñ

Completing the Configuration

After you have created the SplitSite configuration, change the quorum service port, if necessary. Then, enable quorum within the everRun Availability Console. Finally, verify the configuration and (re-)join VMs.

Prerequisite: Plan and create a SplitSite configuration by first reading Creating a SplitSite Configuration and following its instructions, if you have not already done so.

Note: The port configured for quorum service on the quorum computer and the port configured for the quorum server within the everRun Availability Console must be the same port numbers. If you change the quorum service ports on the quroum computer, you must change the quorum service ports on all everRun systems (using the everRun Availability Console) that connect to that quorum computer so that both the quorum computer and the everRun systems use the same port numbers. See Configuring the Quorum Server Within the everRun Availability Console.

Configuring the Quorum Service Port

By default, the quorum service listens on UDP port 4557.

In most cases, you do not need to change the default port. However, you can change the port, if the network configuration requires you to:

To change the port number on the quorum server

- 1. Log on to the quorum computer using an account with administrative privileges.
- 2. Open a command window in administrative mode.
- 3. Stop the quorum service by typing:

net stop sraqserver

4. Change the port by typing (replacing *nnnn* with the new port number):

sraqserver -install nnnn

5. Restart the quorum service by typing:

net start sraqserver

Verifying the Quorum Service Port

If you need to verify the quorum service port, check this Windows registry key:
HKEY_LOCAL_ MACHINE\SYSTEM\CurrentControlSet\Services\SraQserver\Parameters\ QSServerPortForReceive

Configuring the Quorum Server Within the everRun Availability Console

Once the quorum service is running, you should enable the quorum service within the everRun Availability Console. You can also remove a quorum server.

To enable the quorum service:

- 1. Login to the everRun Availability Console with an account that has administrative privileges.
- 2. Click **Preferences** in the left-hand navigation panel, to open the **Preferences** page.
- 3. Click Quorum Servers. The quorum configuration page opens.
- 4. Click Add Quorum Server at the left side of the page.
- 5. In the Add Preferred Quorum Server dialog box, enter the following values (if a preferred quorum server already exists, the Add Alternate Quorum Server dialog box appears):
 - DNS or IP Address—Type the fully-qualified DNS host name or IP address for the preferred quorum server.
 - Port The default port is 4557. Type a port number if you need a port that is different from the default. You need to type only one port number. The quorum service will open the port number for Port and the next port (for example, 4557 and 4558)



Note: The port number must match the port that the quorum service is listening on. (If necessary, you can change the port on the quorum server.)

Click **Save** to save the values.

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

Changes to the quorum configuration do not effect running VMs. You must stop and restart any running VMs after changing the quorum configuration.

To remove a quorum server

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

- 1. Navigate to the **Preferences** page of the everRun Availability Console.
- 2. Click Quorum Servers.

П

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

Note: If a VM is using the quorum server that you are removing, you must reboot the VM so that it no longer recognizes the quorum server, which allows the removal process to finish. The VM will downgrade to simplex mode until it is restarted with no quorum servers configured.

Verify the Configuration and (Re-)Join VMs

Verify the configuration and (re-)join VMs. Follow the appropriate steps in From Initial Installation to Completing the SplitSite Configuration.

Understanding Quorum's Effect on System Behavior

A quorum server in a SplitSite system changes the system's availability and recovery behavior. To understand the quorum's effect on system behavior, you first need to understand the behavior of a system that does not have a quorum server.

Prerequisite: Plan and create a SplitSite configuration by first reading Creating a SplitSite Configuration and following its instructions, if you have not already done so.

An everRun system is designed to provide high availability for one or more guest VMs, which allows the VMs to continue to run even during failures that would otherwise create application downtime. The everRun system can continue to run guest VMs even with, for example, the loss of a single network connection, a hard disk, or even an entire computer.

However, if more catastrophic faults occur (for example, the loss of all possible network paths), the everRun system attempts to determine the overall state of the total system. The system then takes the actions necessary to protect the integrity of the guest VMs.

The following examples illustrate the system's process during a catastrophic fault.

Example 1: A System Without a Quorum Server Experiences a Split-brain Condition

In this SplitSite example, the everRun system includes node0 and node1, but does not include a quorum server. Operation is normal; no faults are currently detected. The two nodes communicate their state and availability over the A-Link connections, as they do during normal (faultess) operation. The following illustration shows normal connections.



A Catastrophic Fault

A careless fork-truck operator crashes through the wall, severing all of the network connections (both business and A-Links), while leaving the power available and the system running. The following illustration shows the fault condition.



Fault Handling

The two nodes handle the fault, as follows:

- Node0—The AX on node0 detects the loss of both A-Links as well as all other network paths. Since the
 node0 AX can no longer detect the presence of its partner, the node0 AX becomes active and runs the
 guest VM. The application inside the guest VM continues to run, perhaps in a limited capacity due to
 the loss of the network.
- Node1—The AX on node1 also detects the loss of both A-Links, but ibiz0 remains available. As its partner does not respond to messages on ibiz0, the node1 AX is now active. The application inside the guest VM continues to run, perhaps not noticing any problems with the system.

From the perspective of an application client or an external observer, the guest VMs are both active and generate network messages with the same return address. Both guest VMs generate data and see different amounts of communication faults. The states of the guest VMs becomes more divergent over time.

Recovery and Repair

After some time, network connectivity is restored: the wall is repaired and the network cables are replaced. When each AX of the AX pair realizes that its partner is back online, the AX pair with the fault handler rules choose the AX that continues as active. The choice is unpredictable and does not include any consideration for which node's performance was more accurate during the split-brain condition.

The data that was generated from the (now) Standby node is overwritten by the resynchronization of the Active node, and thus the data on the (now) Standby node is lost forever.

After a split-brain condition, the system requires several minutes to resynchronize, depending on how much disk activity needs to be sent to the standby node. If several guest VMs are running with different Active nodes, synchronization traffic may occur in both directions.

Note: In some cases, the everRun system may not be able to determine the best way to proceed after a catastrophic fault. In this case, a person needs to recover the system. The recommended recovery method is to use the everRun Availability Console to shut down and reboot one node while the other node continues to run. This method typically forces the running node to become Primary and the AX on that node becomes Active. After the running node becomes Primary, a person can power on the other node. Do not shut down either node if resynchronization is already in progress.

Example 2: A SplitSite System With a Quorum Server Avoids a Split-brain Condition

In this SplitSite example, the everRun system includes node0 and node1 with connections identical to those of the system in Example 1. In addition, the system in Example 2 includes a quorum server. The following illustration shows these connections.



A Catastrophic Fault

П

That careless fork-truck operator crashes through the wall again, severing all of the network connections while leaving the power available and the system running. The following illustration shows the fault condition.



Fault Handling

The two nodes handle the fault, as follows:

- Node0—The AX on node0 detects the loss of both A-Links as well as all other network paths. Since the
 node0 AX can no longer detect the presence of its partner, the node0 AX attempts to contact the
 quorum server. In this case, the quorum server is also unavailable. Therefore, the node0 AX decides to
 shut down. The shutdown is not a graceful Windows shutdown, but is, instead, an abrupt stop, which
 causes the application inside the guest VM to stop.
- Node1—The AX on node1 also detects the loss of both A-Links, but ibiz0 remains available. The node1
 AX tries to contact the quorum server, which responds, so the node1 AX remains active. The applic ation inside the guest VM runs, perhaps not noticing any problems with the system.

Note: If the node1 AX was not previously active and the guest VM is an HA VM, the guest VM on node1 might need to boot from node1's hard drive. In this case, the application experiences a brief period of downtime while the guest VM boots. (FT VMs continue to run.)

From the perspective of an application client or an external observer, the guest VM on node1 remains active and generates data while the VM on node0 is shut down. No split-brain condition exists.

Recovery and Repair

П

After some time, network connectivity is restored: the wall is repaired and the network cables are replaced. When the node1 AX realizes that its partner is back online, the node0 AX becomes Standby. Because node0 was not previously running, data synchronization begins from node1 to node0.

Since a split-brain condition did not occur, no data is lost.

The system requires a few minutes to resynchronize, depending on how much disk activity needs to be sent to the standby node.

Example 2, Modified: The Quorum Server Is Unreachable During the Catastrophic Fault

In a SplitSite system with a quorum server, the quorum server may be offline or otherwise unreachable when the catastrophic fault severs all of the network connections, though the power remains available and the system is still running. The following illustration shows a system in this situation with a quorum server that is off-line.



The fault handling is similar to Example 2 fault handling, with one important difference for node1:

The node1 AX also detects the loss of both A-Links, but ibiz0 remains available. The node1 AX tries to contact the quorum server, but the communication fails. The AX terminates the guest VM.

In this case, the guest VM is shut down on both node0 and node1, preventing split-brain from occurring. The tradeoff is that the guest VM is unavailable until the connection to either node0 or to the quorum server is restored.

In this case, determine the node that you do not wish to operate and power it down. Then, forcibly boot the node that you wish to operate, and then forcibly boot the VM. For information on shutting down a VM and then starting it, see Managing the Operation of a Virtual Machine.)

Example 2, Modified: The Quorum Server Is Unreachable With No Catastrophic Fault

In some situations, the quorum server might be unreachable even without a catastrophic physical failure. One example is when the quorum computer is rebooted for routine maintenance such as applying an OS patch. In these situations, the AX detects that the quorum service is not responding and so the AX suspends synchronization traffic until the connection to the quorum server is restored. The guest VM continues to run on the node that was active when the connection was lost. However, the guest VM does not move to the standby node because additional faults may occur. After the quorum service is restored, the AX resumes synchronization and normal fault handling, as long as the connection to the quorum server is maintained.

Recovering From a Power Failure

If you restart the system after a power loss or a system shutdown, the everRun system waits indefinitely for its partner to boot and respond before the system starts any guest VMs. If the AX that was previously active can contact the quorum server, the AX starts the guest VM immediately without waiting for the partner node to boot. If the AX that was previously standby boots first, it waits for its partner node.

If the system receives a response from either the partner node or the quorum server, normal operation resumes and the VM will start, subject to the same fault handler rules that apply in other cases.

If the system does not receive a response from the quorum server, or if the system does not have a quorum server, then a person must forcibly boot a guest VM, which overrides any decisions made by the AX or the fault handler. You must ensure that two people do not forcibly boot the same guest VM on node0 and node1. Doing so inadvertently causes a split-brain condition.

Accessing Knowledge Base Articles

The **Stratus Customer Service Portal** provides a searchable **Knowledge Base** with technical articles about all Stratus products, including 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

12

Chapter 12: 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.

For information on using the snmptable command to obtain information about the system, specifically information about alerts, audit logs, nodes, VMs, and volumes, see Obtaining System Information with snmptable.

You can download a copy of the MIB file from the **Drivers and Tools** section of the **Downloads** page at https://www.stratus.com/services-support/downloads/?tab=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
```

Field	Description
everRunAlertIndex	The alert number.
everRunAlertSeverity	The alert severity (see everRunAlertSeverityNum for numerical value). Values are: clear \checkmark informational $@$ minor \bigstar major \diamondsuit serious \bigcirc critical \bigcirc
everRunAlertType	<pre>The type of alert. Examples are: node_singleSystemDisk Node Maintenance The Unit is not well balanced</pre>
everRunAlertSource	 The source of the alert. Examples are: node0 or node1 everRun system network name network host name
everRunAlertDateTime	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).
everRun AlertCallHomeSent	If true, Call Home was sent; if false, it was not sent

Field	Description	
everRunAlertEAlertSent	If true, e-Alert was sent; if false, it was not sent	
everRun AlertSNMPTrapSent	If true, SNMP trap was sent; if false, it was not sent	
	Information about the alert. Examples are:	
	• Node nodel is in maintenance	
everRunAlertInformation	 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)	
everRunAlertSeverityNum	everRunAlertSeverity number. Values are:	
	0 Clear 🗸	
	1 Informational	
	2 Minor 🔔	
	3 Major 🧾	
	4 Serious	
	5 Critical	

To display audit log information

To display information about audit logs, issue the following command:

```
snmptable -v2c -m+/usr/smd/STRATUS-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, audit or admin).
everRun AuditOriginatingHost	The IP address of the host that originated the audit.
everRunAuditAction	A description of the action being audited. Examples are: • "Login user \"audit" • "Start virtual machine \"manager1" • "Remove all cleared alert"

To display node information

To display node information, issue the following command:

```
snmptable -v2c -m+/usr/smd/STRATUS-EVERRUN-MIB.txt -c public
localhost everRunNodeTable
```

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:034).	
everRunNodeDisplayName	The node name, node0 or node1.	
everRunNodeIsPrimary	If true, the node is primary. If false, the node is sec- ondary.	
	Node state is:	
	0 Normal (🖍)	
	1 Warning (1)	
everRunNodeStateNum	2 Busy (🏹)	
	3 Broken (🗱)	
	4 Maintenance (
	Node activity is:	
	0 Imaging	
	1 Booting	
	2 Running	
	3 Stopping	
	4 Rebooting	
everRunNodeActivityNum	5 Powered off	
	6 Failed	
	7 Firmware updating	
	8 Lost	
	9 Exiled	
	10 Unreachable	

Field	Description	
	11Proto (initializing)12Evacuating	

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
```

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).	
everRunVMStateNum	VM state is: 0 Normal (**) 1 Warning (▲) 2 Busy or synchronizing (♥) 3 Broken or blacklisted (★)	
everRunVMActivityNum	VM activity is: 0 Installing 1 Booting	

Field	Description	
	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
```

Field	Description
everRunVolumeIndex	An incrementing number (1, 2, etc.) to indicate the volume whose information is displayed.
everRunVolumeId	The volume ID (for example, volume:0588).
everRunVolumeDisplayName	The volume name (for example, root).
everRunVolumeSyncPercentage	The percentage of the volume that is synchronized.

Field	Description
everRunVolumeStorageGroup	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	Volume state is: 0 Normal (*) 1 Warning (*) 2 Busy or synchronizing (*) 3 Broken (*)