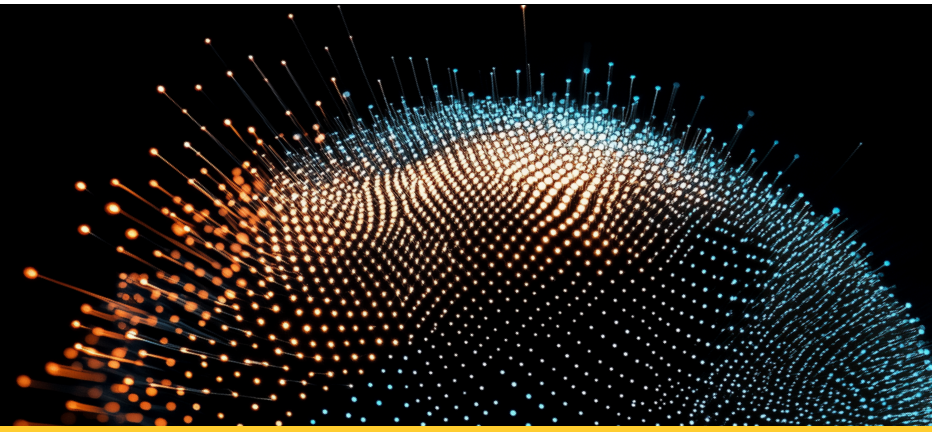
PENGUIN SOLUTIONS | Stratus everRun®



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

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

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

This document is protected by copyright. All rights are reserved. Penguin Solutions 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.

Ubuntu and Canonical are registered trademarks of Canonical Ltd.

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 8.1.1.0

Publication Date: Thursday, October 29, 2026

Penguin Solutions

5 Mill and Main Place, Suite 500

Maynard, Massachusetts 01754-2660

© 2026 Penguin Solutions (PENG) Ireland Limited. All rights reserved.

Table of Contents

Part 1: everRun User's Guide **21**

Chapter 1: Introduction to everRun Systems **23**

everRun System Overview 23

everRun System Description 24

Physical Machines and Virtual Machines 24

Administrative Operations 25

Alerts 25

Remote Support 26

Lights-Out Management 26

Third-party Management Tools 27

Modes of Operation 28

High Availability Operation 28

Fault Tolerant Operation 29

Unprotected Operation 30

SplitSite Configurations 31

SplitSite and Quorum Service 31

Quorum Servers 32

everRun Storage Architecture 32

Logical Disks and Physical Disks 33

Storage Groups 33

Sizing Volume Containers 34

Network Architecture 35

A-Link and Private Networks 35

Business and Management Networks 36

Network Segmentation Fault Detection and Remediation 37

System Usage Restrictions 38

QEMU 38

Accessing the Host Operating System 38

Chapter 2: Getting Started **41**

Planning 41

System Requirements Overview 42

System Hardware 42

Supported Servers 42

RAM ... 42
Disk Space ... 42
Network ... 42
IP Addresses ... 43
Ports ... 43
Storage Requirements ... 44
Memory Requirements ... 46
General Network Requirements and Configurations ... 46
Requirements ... 46
Recommended Configurations ... 47
Business and Management Network Requirements ... 48
A-Link and Private Network Requirements ... 50
everRun Availability Console Requirements ... 50
Compatible Internet Browsers ... 51
Power Requirements and Considerations ... 51
Software Installation ... 51
Site and System Preparation ... 51
Connecting Power ... 52
UPS (Optional) ... 52
Obtaining everRun Software ... 54
Obtaining the ISO Image ... 54
Final Step ... 54
Creating Bootable USB Media ... 54
Configuring Settings in the Firmware Setup Utility ... 57
Required Settings ... 58
Recommended Settings ... 58
Installing everRun Software ... 59
Connecting Ethernet Cables ... 59
Installation Options ... 61
Installing Software on the First PM ... 62
Mapping Your Keyboard ... 68
Recording the Management IP Address ... 69
Installing Software on the Second PM ... 70
Post-Installation Tasks ... 72
Obtaining System IP Information ... 73

Logging On to the everRun Availability Console for the First Time ... 73
Connecting Additional Networks ... 76
Chapter 3: Using the everRun Availability Console ... 79
everRun Availability Console ... 80
Logging On to the everRun Availability Console ... 81
Editing Your User Information ... 83
Dashboard Page ... 84
Resolving Outstanding Alerts on the Dashboard ... 84
System Page ... 85
Rebooting the System ... 86
Shutting Down the System ... 87
Preferences Page ... 88
Specifying Owner Information ... 91
Managing the Product License ... 92
Managing Software Updates ... 95
Configuring IP Settings ... 97
Configuring Quorum Servers ... 99
Configuring Date and Time ... 100
Configuring System Resources ... 102
Configuring the Mail Server ... 102
Configuring Users and Groups ... 104
Managing Local User Accounts ... 105
Managing Domain User Accounts ... 107
Configuring Active Directory ... 108
Configuring the Migration Policy ... 110
Configuring Secure Connections ... 111
Configuring the Console Inactivity Logout ... 116
Configuring the Host Inactivity Logout ... 117
Disabling and Enabling Snapshots ... 117
Configuring VM Devices ... 118
Managing IPtables ... 119
Configuring the Login Banner ... 124
Configuring e-Alerts ... 125
Configuring SNMP Settings ... 126
Configuring Remote Support Settings ... 131

Configuring Internet Proxy Settings ... 133
Alerts History Page ... 134
Audit Logs Page ... 135
Support Logs Page ... 135
Creating a Diagnostic File ... 136
Uploading a Diagnostic File to Customer Support ... 137
Deleting a Diagnostic File ... 138
Physical Machines Page ... 139
Physical Machine Actions ... 140
Physical Machine States and Activities ... 141
Virtual Machines Page ... 142
Virtual Machine Actions ... 144
Virtual Machine States and Activities ... 147
Snapshots Page ... 149
Volumes Page ... 150
Storage Groups Page ... 151
Networks Page ... 152
Fixing a Network Connection ... 153
Setting the MTU ... 154
Virtual CDs Page ... 154
Upgrade Kits Page ... 155
Creating a USB Medium with System Software ... 156
Chapter 4: Upgrading everRun Software ... 159
Upgrading everRun Software Using an Upgrade Kit ... 159
Upgrading from everRun Release 7.x to 8.x or Higher ... 163
Migrating Virtual Machines to 512e Storage ... 173
Chapter 5: Managing Logical Disks ... 177
Logical Disk Management ... 177
Responding to a Failed Logical Disk ... 178
Activating a New Logical Disk ... 178
Creating a New Storage Group ... 179
Deleting a Storage Group ... 180
Assigning a Logical Disk to a Storage Group ... 181
Chapter 6: Managing Physical Machines ... 183
Maintenance Mode ... 183

Rebooting a Physical Machine ... 185
Shutting Down a Physical Machine ... 186
Load Balancing ... 187
Modes of Operation ... 187
Troubleshooting Physical Machines ... 188
Recovering a Failed Physical Machine ... 188
Chapter 7: Managing Virtual Machines ... 195
Planning Virtual Machine Resources ... 196
Planning Virtual Machine vCPUs ... 197
Planning Virtual Machine Memory ... 198
Planning Virtual Machine Storage ... 199
Planning Virtual Machine Networks ... 201
Adding, Migrating, and Importing Virtual Machines ... 202
Creating a New Virtual Machine ... 203
Copying a Virtual Machine ... 209
Backing Up a Running Virtual Machine ... 212
Recovering a Virtual Machine from a Backup ... 216
Migrating a Physical Machine or Virtual Machine to a System ... 220
Importing Virtualization Files ... 230
Importing Virtual Machines from a VCD ... 240
Replacing/Restoring a Virtual Machine from an OVF File ... 243
Exporting a Virtual Machine ... 249
Mounting a USB Device or Network-mounted Folder on the everRun System ... 254
Managing Windows Drive Labels ... 256
Understanding Windows 11 Virtual Machine Differences from Other VMs ... 257
Configuring Windows-based Virtual Machines ... 259
Updating the VirtIO Drivers (Windows-based VMs) ... 259
Creating and Initializing a Disk (Windows-based VMs) ... 262
Installing Applications (Windows-based VMs) ... 263
Installing the QEMU Guest Agent for Application-Consistent Snapshots (Windows-based VMs) ... 264
Configuring Linux-based Virtual Machines ... 266
Creating and Initializing a Disk (Linux-based VMs) ... 267
Installing Applications (Linux-based VMs) ... 267
Installing the QEMU Guest Agent for Application-Consistent Snapshots (Linux-based VMs) ... 268

Managing the Operation of a Virtual Machine 269
Starting a Virtual Machine 270
Shutting Down a Virtual Machine 270
Powering Off a Virtual Machine 272
Opening a Virtual Machine Console Session 273
Renaming a Virtual Machine 277
Removing a Virtual Machine 277
Managing Virtual Machine Resources 278
Reprovisioning Virtual Machine Resources 279
Creating a Volume in a Virtual Machine 282
Attaching a Volume to a Virtual Machine 284
Detaching a Volume from a Virtual Machine 285
Removing a Volume from a Virtual Machine 287
Renaming a Volume on the everRun System 288
Expanding a Volume Container on the everRun System 289
Expanding a Volume on the everRun System 290
Recovering Virtual Machine Resources 292
Enabling and Disabling VM Components 292
Managing Virtual CDs 293
Creating a Virtual CD 294
Inserting a Virtual CD 296
Ejecting a Virtual CD 297
Booting from a Virtual CD 297
Renaming a Virtual CD 298
Downloading a Virtual CD 298
Removing a Virtual CD 299
Managing Snapshots 299
Creating a Snapshot 301
Creating a Virtual Machine from a Snapshot 304
Exporting a Snapshot 306
Removing a Snapshot 310
Advanced Topics (Virtual Machines) 312
Assigning a Specific MAC Address to a Virtual Machine 313
Selecting a Preferred PM for a Virtual Machine 314
Forcing a VM to Boot 314

Changing the Protection Level for a Virtual Machine (HA or FT) 318
Configuring the Boot Sequence for Virtual Machines 319
Resetting MTBF for a Failed Virtual Machine 320
Locating a Dump File in a Virtual Machine 321
Attaching a USB Device to a Virtual Machine 322
Chapter 8: Maintaining Physical Machines 325
Physical Machine Hardware Maintenance Restrictions 326
Adding or Replacing Hot-Swappable Components 327
Adding or Replacing Components That Are Not Hot-Swappable 327
Adding a New NIC 329
Replacing Physical Machines, Motherboards, NICs, or RAID Controllers 330
Upgrading Both Physical Machines In a Running System 341
Part 2: Supporting Documents 343
Chapter 9: everRun Release 8.1.1.0 Release Notes 345
New Features and Enhancements 345
New Features in Previous Releases 346
Features No Longer Supported 346
Legacy Upgrade Paths No Longer Supported 346
Bug Fixes 347
Upgrade Now Verifies that Priv0 Is Isolated 347
Restriction on Number of Virtual CPUs in Windows 11 Environments 347
VM Creation Fails When Using a Remote ISO (NFS or CIFS) If the File Path Contains Spaces 347
Upgrade Process Now Validates SSD Life Remaining 347
Cannot Attach USB Device to Unprotected VM 348
Logical Disks Become Foreign on Node After Power Interruption 348
Bugs Fixed in everRun Release 8.1.0.2 348
CVE Fixes 348
Important Considerations 348
Upgrading to everRun Release 8.1.1.0 349
Dell PowerEdge Systems and UEFI0074 Prompt During Installation 349
Requirements for Upgrading Windows 10 Virtual Machine Guest Systems to Windows 11 349
REST API Restrictions for Importing a VM 350
SNMP Disabled by Default on everRun Systems 350
Tested Guest Operating Systems 350

Known Issues 351
Fault-Tolerant VM Becomes Stuck After Rapid USB Detach and Re-Attach 351
Rolling Reboot Hangs Indefinitely When a VM Has an Attached USB Storage Device 351
Backup of Stopped Virtual Machine Triggers Unnecessary Warning 352
Duplicate VirtIO VCD Created After Power Cycles 352
Intermittent Snaphot Creation Failure on Windows Server 2022 Virtual Machines 352
Mismatch in OVA File Name Validation Between everRun Availability Console and REST API 352
VM Import Succeeds With Invalid Destination Parameter 353
Swagger UI Accepts Invalid Authorization Tokens Without Immediate Error Message 353
UEFI Debian VMs Fail to Boot After P2V Conversion 353
P2V Conversion Fails for Rocky Linux Version 9.4 353
everRun Legacy BIOS Installation with Virtual Floppy Starts from sdb Disk 353
Unable to Create Application-Consistent Snapshot for Ubuntu 24.04.1 VM 353
DVD/USB Recovery or Replacement Fails on System with 512e, 512n and 4K Logical Disks 354
Guest Performance Issues with Large Guest Volumes 354
Graphical Installation Fails to Start When Installing Debian 11 Guest 354
Removable Media and Migrating a PM or VM Using the P2V Client 354
"The VM name has failed to start" Alert While Running the P2V Client Is Normal 354
Unknown Error in everRun Availability Console When Importing an Unprotected VM 355
AVCLI Import or Restore of VM with Two VHDX Volumes Fails 355
Maximum Resolution of a UEFI VM Console Session 355
Creating VCD Fails When Console Browser Is Microsoft Edge 355
Mapping of Japanese Keyboards 106 and 109 for Console in Firefox May Be Incorrect 355
Cannot Enable SNMP Requests Without Traps 355
System Fails to Generate SNMPv3 Traps When Authentication Is Enabled 356
VMs Running Windows 2016 with the Maximum vCPUs and Memory Will Not Reboot Cleanly 356
Some Browsers Unable to Connect a VNC When Using https 356
Some Browsers Unable to Connect a VNC When Using http 356
Reboot Required When Changing Node IP Address or Netmask Network Settings 356
Thermald “Unsupported CPU model or platform” Messages During Node Startup 357
everRun Installation Fails on First Boot on a System with Virtual Drives 357
everRun Installation Fails on a Previously Installed System with EFI Boot Entries 357
Cannot Attach USB Device to VM on First Attempt 358
Incomplete Messages When Creating or Importing a Windows 11 Guest with the BIOS Boot Type 358

Windows Server 2016 Guest Stuck in Booting State After Installation in Localized German Language ... 358
Accessing Knowledge Base Articles ... 359
Getting Help ... 359
Chapter 10: everRun Command Line Interface Reference ... 361
AVCLI Command Overview ... 361
Prerequisites ... 362
Installing the Linux Client ... 362
Installing the Windows Client ... 363
Using AVCLI ... 364
Executing a Command ... 364
Using AVCLI Help ... 365
Listing All Commands ... 365
Displaying Help for a Specific Command ... 366
AVCLI Error Status ... 366
XML Encapsulated Errors ... 366
Error Checking ... 367
Asynchronous Command Delay ... 367
Output Formatting ... 368
User-Friendly Command Output ... 368
Program-Friendly XML Output ... 369
AVCLI Exceptions ... 372
AVCLI Command Descriptions ... 373
ad-disable ... 379
ad-enable ... 380
ad-info ... 381
ad-join ... 382
ad-remove ... 383
alert-delete ... 384
alert-info ... 385
audit-export ... 386
audit-info ... 387
callhome-disable ... 388
callhome-enable ... 389
callhome-info ... 390

datetime-config 391
diagnostic-create 394
diagnostic-delete 395
diagnostic-extract 396
diagnostic-fetch 397
diagnostic-info 398
dialin-disable 399
dialin-enable 400
dialin-info 401
disk-move-to-group 402
ealert-config 403
ealert-disable 404
ealert-enable 405
ealert-info 406
help 407
image-container-info 408
image-container-resize 411
kit-add 412
kit-delete 413
kit-info 414
kit-qualify 415
kit-upgrade 416
kit-upgrade-cancel 417
license-info 418
license-install 419
local-group-add 420
local-group-delete 421
local-group-edit 422
local-group-info 423
local-user-add 424
local-user-delete 426
local-user-edit 427
local-user-info 429
localvm-clear-mtbf 430
mail-server-config 431

mail-server-disable 433
mail-server-enable 434
mail-server-info 435
media-create 436
media-delete 437
media-eject 438
media-import 439
media-info 440
media-insert 441
network-change-mtu 442
network-change-role 444
network-info 445
node-add 447
node-cancel 448
node-config-prp 449
node-delete 450
node-delete-prp 451
node-info 452
node-reboot 453
node-recover 454
node-shutdown 455
node-workoff 456
node-workon 457
ntp-config 458
ntp-disable 459
ova-info 460
ovf-info 461
owner-config 462
owner-info 463
pm-clear-mtbf 464
proxy-config 465
proxy-disable 466
proxy-enable 467
proxy-info 468
removable-disk-info 469

snmp-config 470
snmp-disable 472
snmp-info 473
snmp-v3-add-agent-user 474
snmp-v3-add-trap-recipient 477
storage-group-create 480
storage-group-delete 481
storage-group-info 482
storage-group-info-v2 484
storage-info 486
timezone-config 487
timezone-info 488
unit-avoid-bad-node 489
unit-change-ip 490
unit-configure 492
unit-eula-accept 493
unit-eula-reset 494
unit-info 495
unit-shutdown 496
unit-shutdown-cancel 497
unit-shutdown-state 498
unit-synced 499
vm-attach-usb-storage 500
vm-ax-disable 502
vm-ax-enable 503
vm-boot-attributes 504
vm-cd-boot 505
vm-copy 506
vm-create 510
vm-create-from-snapshot 515
vm-delete 517
vm-device-config-info 518
vm-export 519
vm-import 521
vm-info 524

vm-media-insert-disable ... 525
vm-media-insert-enable ... 526
vm-network-disable ... 527
vm-network-enable ... 528
vm-poweroff ... 529
vm-poweron ... 530
vm-reprovision ... 531
vm-restore ... 535
vm-shutdown ... 538
vm-snapshot-create ... 539
vm-snapshot-create-disable ... 541
vm-snapshot-create-enable ... 542
vm-snapshot-delete ... 543
vm-snapshot-export ... 544
vm-snapshot-info ... 546
vm-unlock ... 547
vm-usb-attach-disable ... 548
vm-usb-attach-enable ... 549
vm-volume-disable ... 550
vm-volume-enable ... 551
volume-info ... 552
volume-resize ... 553
Chapter 11: System Reference Information ... 555
Tested Guest Operating Systems ... 555
Physical Machine System Requirements ... 558
Important Physical Machine and Virtual Machine Considerations ... 562
Virtual Machine Recommendations and Limits ... 562
Recommended Number of CPU Cores ... 562
Virtual Machine Limits ... 563
Combined Virtual Machine Maximums ... 564
Important Considerations ... 564
Creating a SplitSite Configuration ... 565
Creating the Configuration ... 570
A Typical everRun System ... 570
A SplitSite Configuration With a Quorum Server ... 571

SplitSite VLAN Requirements ... 572
From Initial Installation to Completing the SplitSite Configuration ... 572
Meeting Network Requirements ... 574
Locating and Creating the Quorum Server ... 576
Locating the Quorum Computer ... 576
Adding an Alternate Quorum Computer ... 577
Quorum Computer Requirements ... 578
Downloading and Installing the Quorum Service Software ... 578
Completing the Configuration ... 579
Configuring the Quorum Service Port ... 579
Verifying the Quorum Service Port ... 580
Configuring the Quorum Server Within the everRun Availability Console ... 580
Verify the Configuration and (Re-)Join VMs ... 581
Understanding Quorum's Effect on System Behavior ... 581
Example 1: A System Without a Quorum Server Experiences a Split-brain Condition ... 582
A Catastrophic Fault ... 582
Fault Handling ... 583
Recovery and Repair ... 583
Example 2: A SplitSite System With a Quorum Server Avoids a Split-brain Condition ... 584
A Catastrophic Fault ... 584
Fault Handling ... 585
Recovery and Repair ... 585
Example 2, Modified: The Quorum Server Is Unreachable During the Catastrophic Fault ... 586
Example 2, Modified: The Quorum Server Is Unreachable With No Catastrophic Fault ... 587
Recovering From a Power Failure ... 587
Accessing Knowledge Base Articles ... 587
Fixed CVEs ... 588
CVEs Fixed in everRun Release 8.1.1.0 ... 588
CVEs Fixed in everRun Release 8.1.0.2 ... 608
CVEs Fixed in everRun Release 8.1.0.1 ... 609
CVEs Fixed in everRun Release 8.1.0.0 ... 610
CVEs Fixed in everRun Release 8.0.0.0 ... 616
CVEs Fixed in everRun Release 7.9.4.1 ... 617
CVEs Fixed in everRun Release 7.9.4.0 ... 617
CVEs Fixed in everRun Release 7.9.3.1 ... 618

CVEs Fixed in everRun Release 7.9.3.0 618
CVEs Fixed in everRun Release 7.9.2.0 619
CVEs Fixed in everRun Release 7.9.1.1 620
CVEs Fixed in everRun Release 7.9.1.0 620
CVEs Fixed in everRun Release 7.9.0.0 621
CVEs Fixed in everRun Release 7.8.0.0 627
CVEs Fixed in everRun Release 7.7.0.0 631
CVEs Fixed in everRun Release 7.6.1.0 636
CVEs Fixed in everRun Release 7.6.0.0 639
REST API 640
Accessing the REST API documentation 641
Logging In for REST API Testing 642
Extending a Login Session with the Refresh Token 643
Logging Out of the REST API 644
List of REST API Methods 644
Chapter 12: Security 647
Security Hardening 647
Security Guidelines 648
Ports and Protocols 649
Network Segmentation 649
IP Tables/Firewall 649
User Account Creation 650
Password Creation 650
Least Privilege 651
Active Directory 651
BIOS 651
Ports 652
Time Synchronization 652
Secure Connections 652
Updating SSL Certificate 653
SNMP Configurations 653
Backups 654
SplitSite Configuration 654
Auditing 654
Login Banner Notice 655

Upgrades 655
Physical Security 656
Advanced Security Guidelines 656
Password Quality Recommendations 656
Concurrent User Management 657
Antivirus 657
SSH Access Restrictions 658
Best Practices and Standards of Standards Organizations 659
Chapter 13: SNMP 663
Obtaining System Information with snmptable 663

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 (Ubuntu)
- Contain the same data, memory and storage (synchronized via direct Ethernet links between the two PMs)
- Support virtual machines running tested 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.

everRun systems provide a secure out-of-box experience. You also have the option of implementing additional security-related configuration settings. For information, see Security.

Related Topics

System Requirements Overview

Tested 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

Physical Machines Page

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.

 Informational

 Normal or OK state

 Minor, warning, or inconsistent state

 Moderate state

 Broken, failed, or severe state

Remote Support

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

- **Support Configuration**–Configure settings to allow remote support access of your system by your authorized Penguin Solutions service representative and to enable your system to send health and status notifications to your authorized Penguin Solutions 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 (Ubuntu) should run on everRun systems. Possible exceptions are tools that manage/monitor the KVM-based virtualization. To manage/monitor everRun virtualization, use the integrated everRun management tools.
- Before deploying your everRun system, ensure that you verify that it operates properly with the management tools installed and operational.
- 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 Penguin Solutions service representative.

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

Related Topics

Getting Started

System Reference Information

Modes of Operation

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

- High Availability Operation
- Fault Tolerant Operation
- Unprotected Operation

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

It is important to configure 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.

In unprotected operation, a VM runs on only one PM; if that PM shuts down or fails, the VM will not fail over to the second PM. Use unprotected operation only if a VM has specific requirements to do so, for example, if software licensing restrictions prevent your applications from being installed or running at the same time on both PMs for redundancy.

High Availability Operation

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

In High Availability (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.

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

Modes of Operation

Virtual Machines Page

Using the everRun Availability Console

Fault Tolerant Operation

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

Use Fault Tolerant (FT) operation for applications that need the highest levels of availability. In FT operation, an application continues to run without downtime during a fault. 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

Modes of Operation

Virtual Machines Page

Using the everRun Availability Console

Unprotected Operation

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

In unprotected operation, a virtual machine (VM) runs on only one physical machine (PM, or node); if that PM goes offline for any reason, the VM stops running. Use unprotected operation only if a VM has specific requirements to do so, for example, if software licensing restrictions prevent your applications from being installed or running at the same time on both PMs for redundancy.



Caution: If you set the **Protection Level** of a VM to **NONE**, the VM is unprotected. If the preferred node shuts down or fails, the VM and its applications will stop. The VM will not fail over to the second node.

Unprotected operation has the following configuration requirements and behavior:

- You can configure a VM for unprotected operation only when you create or import the VM by setting the **Protection Level** to **NONE** and selecting a **Preferred node** on which the VM will run.
- The option to set the protection level is available only on dual-node systems. On single-node systems, the VMs automatically run in unprotected mode.
- After you initially create or import the VM, you cannot modify its protection level or switch the preferred node (in the **Reconfigure Virtual Machine** wizard, for example). If you need to convert the VM to an HA- or FT-protected VM, you can do so only by exporting and re-importing the VM.
- If you put the preferred node into maintenance mode (**Work On**) or upgrade the everRun system software, the system will automatically shut down any unprotected VMs.
- When the preferred node powers back on or comes back online from maintenance mode, the VM either starts or remains stopped depending on its auto start setting. For information about modifying the auto start setting, see Starting a Virtual Machine.
- If you need to replace the preferred node that contains an unprotected VM, you must shut down the unprotected VM, export the VM, and delete the VM before replacing the node; otherwise, you will be unable to boot the VM on the replacement node. After replacing the node, you can import the VM and configure it as an unprotected VM again.
- You cannot copy unprotected VMs. Also, snapshots are not available for unprotected VMs.

- You cannot back up an unprotected VM as described in Backing Up a Running Virtual Machine, but you can recover a VM to an unprotected VM as described in Recovering a Virtual Machine from a Backup.
- An unprotected VM cannot be load balanced. It runs only on the preferred PM.

Related Topics

Modes of Operation

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, it is important to 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. It is important to use 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 everRun software detects new logical disks and logical disk failures. You manage logical disks using the everRun Availability Console. For information, see Managing Logical Disks.

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

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

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

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 (Ubuntu) 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 (Ubuntu) host operating system of the everRun system from any source other than Stratus. Use only the 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.



Note: The default password for the `root` account is **KeepRunning**. To secure the system, change the `root` password on each PM as soon as possible. To change the password, issue the `passwd` command on each PM.

For information about using third-party management tools in the host operating system, 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**.
4. Click **Physical Machines** in the left-hand navigation panel to open the **Physical Machines** page.
5. Record which PM is the primary node for the system, displayed as **node*n* (primary)**. In most cases, log on to the IP address of the primary node to ensure that administrative commands will work properly.

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

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

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

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

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

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

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

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

2

Chapter 2: Getting Started

The following topics describe the 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 Tested Guest Operating Systems.

System Hardware

Supported Servers

Stratus everRun software will run on any systems listed in the Ubuntu® Certified hardware list that support Ubuntu 22.04 LTS 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 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* (KB0012595). 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.
- It is important that your system use a storage RAID controller.
    - If the system has a single logical disk, configure the RAID controller so that logical disks presented to the host are backed by redundant physical drives.
    - RAID controllers should 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 512n (Standard) format, 512e format, and Advanced 4K Native format with sectors, as follows:

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

For best performance, use disks with the 4K sector size. 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 or compatible disk types:
    - A storage group with the 512n disk type supports logical disks with the 512n or 512e disk type.
    - A storage group with the 512e disk type supports logical disks only with the 512e disk type.
    - A storage group with the 4K disk type supports logical disks only with the 4K disk type.

For example, you cannot add a disk with the 4K logical disk disk type to a storage group with the 512n or 512e disk type.

- The disk type of the **Initial Storage Group** is automatically defined by the disk type of the boot disk in the first PM on which you install the everRun software: the 4K disk type for a 4K boot disk, or the 512n disk type for a 512n or 512e boot disk. The boot disk in the second PM must be the same disk type. You cannot change the disk type of the **Initial Storage Group** after the software installation.
- You can set the disk type for other storage groups only when you create each storage group and select its **Disk Type** (as described in Creating a New Storage Group). You cannot change the disk type of an existing storage group; you would need to create a new storage group and select the new disk type.
- Because the disk type 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 512n or 512e disk type provides the sector size of 512B for its VM volumes.
  - A storage group with the 4K disk type provides a sector size of either 4K or 512B, selectable for each of its VM volumes.
  - If a storage group with the 512e or 4K disk type provides a volume with the 512B sector size, it is presented to the VM as a volume with the 512e disk type.

  Note that the boot volume for each VM must be 512B, regardless of the storage group's disk type. 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.
- If a system contains a mix of Serial Attached SCSI (SAS) and Non-Volatile Memory Express (NVMe) devices, the everRun software prevents any NVMe devices from being included in the **Initial Storage Group**, because not all BIOS and UEFI systems allow NVMe devices to be bootable devices. However, if a system contains only NVMe devices, the **Initial Storage Group** may contain NVMe 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

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

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.



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 one or a maximum of two 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 everRun A-Link Networks* (KB0013457). 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.

Compatible Browsers	Release
Microsoft Edge	42.17134 or greater
Mozilla® Firefox®	65.0 or greater
Google® Chrome™	73.0 or greater

Power Requirements and Considerations

To ensure maximum availability, everRun's fault-tolerant (FT) software must 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

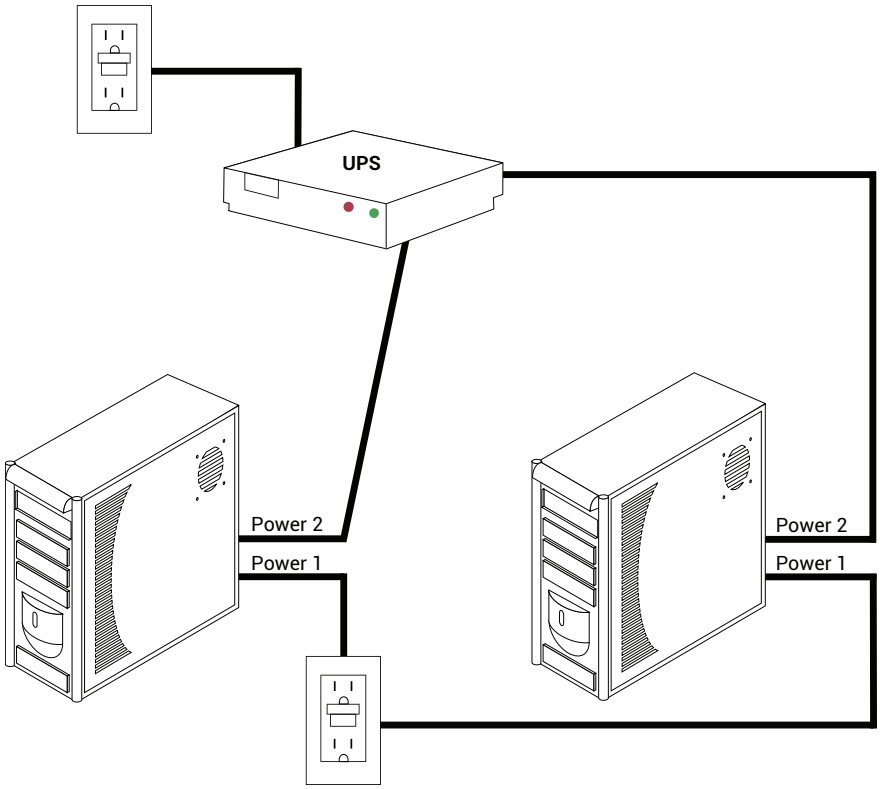
To connect power, configure a dual-node everRun system with redundant power supplies connected to separate sources. You can optionally use uninterruptible power supplies (UPS), as shown below.

After connecting power, return to Installing everRun Software.

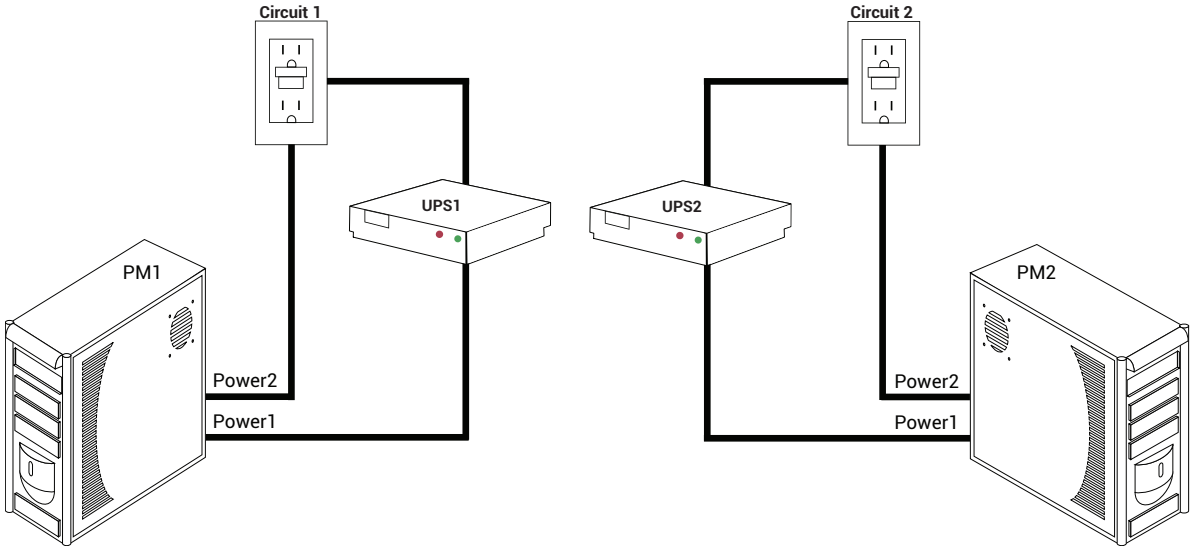
UPS (Optional)

Note: 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

1. From any computer connected to the Internet, go to the **Downloads** page at https://www.penguinsolutions.com/services-support/downloads/?tab=everrun.
2. To download the everRun software ISO image (**everRun_install-8.*x.x.x-xxx*.iso**), under **Product Downloads**, click **everRun 8.*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/sd*x* bs=*blocksize***, where ***path_to_image*** is the full path to the ISO image file that you saved, **sd*x*** 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-8.0.0.0-
   nnn.isoof=/dev/sdc bs=8K
   ```

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

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

Creating a bootable USB medium on a Windows-based system

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

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

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

1. Download and save the everRun software ISO image (see Obtaining everRun Software) to a Windows-based system, if you have not already done so.
2. Ensure the integrity of the ISO image using a tool for verifying MD5 checksums on Windows-based systems. For example, use the MD5 checksum hash function. To do so, 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.
3. Download and install the **Rufus** utility at http://rufus.akeo.ie/. Be sure to click the **Download** link that is about halfway down the web page (below **Last updated**); do not click advertisements, including **Download** links that appear in advertisements.
4. Insert a USB medium such as a thumbdrive into the system
5. Run the **Rufus** utility, selecting the following options:

Option	Value(s)
Partition scheme and target system type	MBR partition scheme for BIOS and UEFI

Option	Value(s)
File system	FAT32
Cluster size	Specify the system default
Format Option	Create a bootable disk using ISO Image (browse to the xxx.iso image) Create extended label and icon files

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

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

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

Related Topics

Obtaining everRun Software

Software Installation

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 required.

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-Disable Bit Capability	Allows the processor to classify areas in memory where application code can or cannot execute. Set this to Enabled to help prevent malicious code attacks.

Recommended Settings

The following settings are optional but recommended.

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

Installing everRun Software

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

Warning: Installing everRun software erases all hard drives.

To install everRun software for the first time:

1. On a remote management computer, obtain the everRun software. See Obtaining everRun Software
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.
4. After you have finished installing the software on the first PM, perform the installation on the second PM. See Installing Software on the Second PM.
5. The software installation is complete. Now, perform the required post-installation configuration steps. See Post-Installation Tasks.

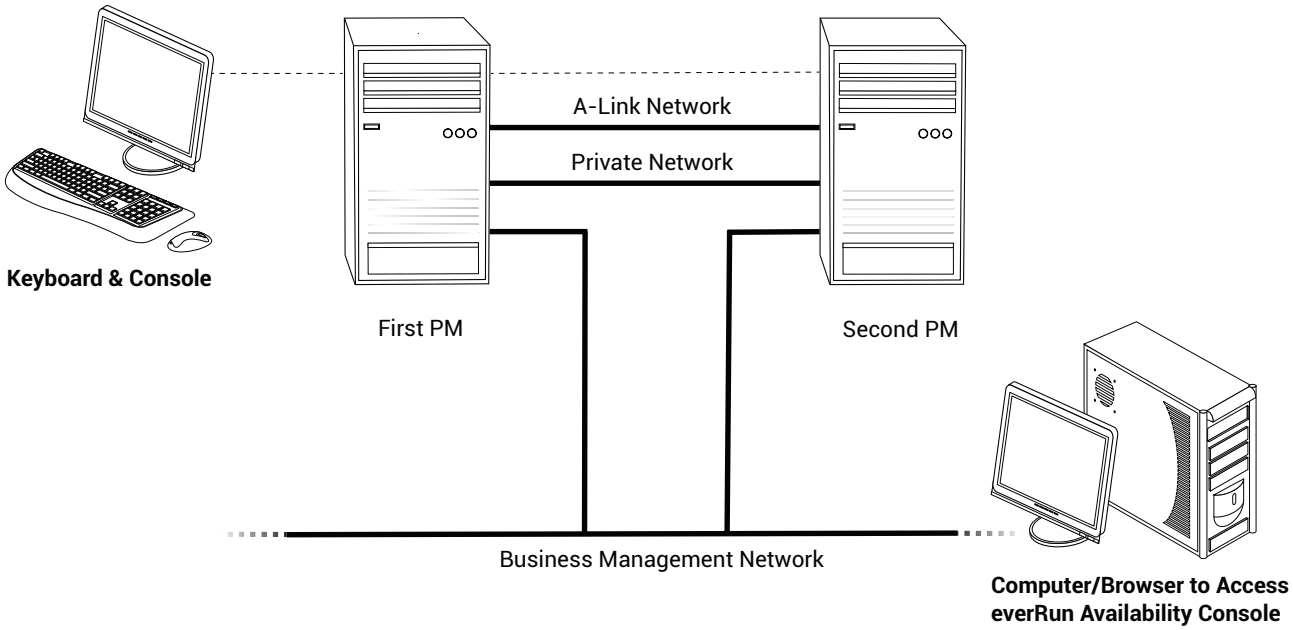
Connecting Ethernet Cables

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

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

On 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, for best results, 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.



See the following recommended Ethernet cable configurations:

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

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

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

  Note: Configure at least one 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 perform the installation.	`Verify medium and Install everRun`	Verifies the installation medium first, then installs the Ubuntu host operating system and the everRun software, and, finally, creates a new system. (You should verify installation medium the first time you use it; note, however, 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 physical 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 physical machine. See Replacing Physical Machines, Motherboards, NICs, or RAID Controllers.
Perform initial install-	`Install`	Deletes all partitions on all connected disks, installs

Task	Option	Description
ation on the first PM	everRun, Create a new system	the Ubuntu host operating system and the everRun software, and creates a new system. 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 Ubuntu host operating system and the everRun software, and attempts to connect to an existing system. (This option is the default.) See Recovering a Failed Physical Machine.
Perform initial installation on the second PM; replace a PM	Replace PM, Join system: Initialize data	Deletes all partitions on all connected disks, installs the Ubuntu host operating system and the everRun software, and attempts to connect to an existing system. See Installing Software on the Second PM and Replacing Physical Machines, Motherboards, NICs, or RAID Controllers.
Boot to rescue mode (UEFI firmware installation only)	Rescue the installed system	Boots to rescue mode.

Installing Software on the First PM

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

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

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

1. Power on the first PM, if it is not already powered on, and either insert the bootable medium or mount the ISO image.

2. As the system powers on, enter the firmware (BIOS or UEFI) setup utility, and configure the required and optional settings described in Configuring Settings in the Firmware Setup Utility.
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

> **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.

   For an initial installation, you should select **Verify medium and Install everRun**, 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**:
> 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**:
>
> 1. If you are not sure of which port to use, use the arrow keys to select one of the ports, and click the **Identify** button. The LED on the selected port will then flash for 30 seconds, allowing you to identify it. Since the LED may also flash due to activity on that network, you should 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 dynamically, 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

Layout	Language
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
   ```

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

To configure your keyboard layout after installation:

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

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

3. Repeat the preceding steps on the second PM.

Related Topics

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 dynamic IP address. (Your network administrator already has this information if the management network has a static IP address.)

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

```
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

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.

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

   Notes:

   1. If you are not sure of which port to use, use the arrow keys to select one of the ports, and click the **Identify** button. The LED on the selected port will then flash for 30 seconds, allowing you to identify it. Since the LED may also flash due to activity on that network, you should 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.

8. 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.

10. 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 dynamically, 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

> **Note**: You must specify an email address for each user account, including **admin**, to enable the forgot password feature. If a user account does not include an email address, and the user clicks the **Forgot Password?** link on the console login page, the system sends an email to **user@example.com**. Managing Local User Accounts describes how to add users as well as how to edit user accounts, including how to add email addresses.

- 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 **Penguin Solutions Customer Service Portal** at https://support.penguinsolutions.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.

   > **Note**: If a security message appears, proceed to the web site. You can add a security exception later, to allow the site to load without the message (see Configuring Secure Connections).

   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 Penguin Solutions service representative, uncheck the box. You can change this setting later (see Configuring Remote Support Settings).

5. Under **SYSTEM IP**, for **Static System IP**, enter the static system IP address that you obtained from your network administrator. (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

   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.

7. You can upload a new license key now or later from the **Product License** page. To do so later, click **Continue**.

   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**. Click **Continue**.

8. For **New Password** in the **ACCOUNT SECURITY** window, type a new password for the user **admin**. Type the password again in **Confirm Password**. The password must conform to the password policy of the system (for information, see Password Policy).

   Notes:

   - You must change the password for **admin** now, for security. You can change it again later, and you should change the default user login name for the **admin** account. You make these changes on the **Users & Groups** page (see Configuring Users and Groups).
   - For additional security, also change the password for **root** in the host operating system of each PM as soon as possible after installation (see Accessing the Host Operating System).

9. 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

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

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

Related Topics

Connecting Ethernet Cables

A-Link and Private Network Requirements

Business and Management Network Requirements

General Network Requirements and Configurations

3

Chapter 3: Using the everRun Availability Console

The everRun Availability Console is a browser-based interface that provides management and monitoring of everRun systems 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:

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

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. See the Upgrade Kits Page.

You can also edit your user information (see Editing Your User Information) and configure users and groups (see Configuring Users and Groups).

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.

Notes:

1. A login session times out after one hour, if unused.
2. The system has a limit of 10 login sessions.
3. Passwords must conform to the Password Policy of the system.
4. You can configure a login banner to provide customized content to the everRun Availability Console login page. See Configuring the Login Banner.

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**.

   If you have forgotten your password, click **Forgot Password?** and the **Reset Password** page appears. Enter the requested information to reset your password.

   > **Note**: Resetting a password requires that you have an email account on the system, with an email address, as configured in your local user account (see Managing Local User Accounts). If you are unable to receive email, you must contact your system administrator, who will request a password reset for you. (The system administrator needs to ask the administrator of the host OS to change the password. The host OS administrator changes the password by using AVCLI commands on the primary node.)

   To reset your password

   > **Note**: To receive email when resetting your password, the Mail Server must be configured. See Configuring the Mail Server.

   a. When the **Reset Password** page appears, enter your **Username** and click **Continue**. An email is sent to the email address listed with your local user account. The email contains a link to a reset password page.

   b. In your email account, open the email with the reset-password link, and click the link. The **Reset Password** page re-appears.

   c. For **New Password** and **Confirm Password**, type your new password. The new password must conform to the Password Policy of the system.

      Click **Continue**.

   d. A page appears, with a message that the reset was successful and that you can log in to the system with your new password. Click **Finish**.

3. Click **LOGIN**.

Password Policy

The password policy of the system requires that your password meet these conditions:

- Its minimum length is 8 characters.
- It must contain both upper- and lower-case characters.
- It cannot be the username.

> **Note**: The interval between login attempts is 500 ms, so, after a login attempt, you must wait at least a half second to log in again.

Related Topics

Logging On to the everRun Availability Console for the First Time

everRun Availability Console

Using the everRun Availability Console

Editing Your User Information

Edit your user information (that is, your user profile) by changing your user name, email address, real name, or password.

To edit your user information

1. Click your user name in the upper right-hand corner of the console.

   The **Edit User** dialog box opens.

2. Enter or modify values for the following:
   - **User Name**
   - **Email Address**
   - **Real Name**
   - **Password**

     > **Note**: Passwords must confirm to the Password Policy of the system.

   - **Confirm Password**
3. Click **Save**. (Or click **Cancel** to cancel the changes.)

Related Topics

everRun Availability Console

Using the everRun Availability Console

Dashboard Page

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

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

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

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

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

Physical Machines Page

System Page

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

Dashboard Page

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

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

everRun Availability Console

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:

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

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

everRun Availability Console

System Page

Using the everRun Availability Console

Preferences Page

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

The following table lists and describes the preferences.

Preference	Description
System	
Owner Information	Allows you to specify and then view the name and contact information for an everRun system administrator. This information is also provided in response to Simple Network Management Protocol (SNMP) requests. See Specifying Owner Information.
Product License	Allows you to view and manage the everRun product license. See

Preference	Description
	Managing the Product License.
Software Updates	Allows you to check the current version of the system software and whether or not a new version is available. If a new version is available, you can download it and read the Release Notes. You can also specify that alerts be sent when an update is available and that an available update be downloaded automatically. See Managing Software Updates.
IP Configuration	Allows you to view and specify the Internet Protocol (IP) address and network 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 specific failures in the everRun environment. See Quorum Servers and Configuring Quorum Servers.
Date & Time	Allows you to view the system time, specify values for Network Time Protocol (NTP) (recommended), or to manually set the time and date on the system. 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.
Mail Server	Allows you to configure the mail server to enable the everRun system to send email when, for example, someone needs to reset a password. See Configuring the Mail Server.
Administrative Tools	
Users & Groups	Allows you to add, modify, or remove user accounts on the everRun system; to enable Active Directory (and then grant to it), and to select a user and view the time when the user's password was last updated. An administrator can also use the page to force a selected user to change the user's

Preference	Description
	password on the next login. See Configuring Users and Groups.
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 Configuring Secure Connections.
Console Inactivity Logout	Allows you to enable the everRun Availability Console inactivity logout and to change the timeout. See Configuring the Console Inactivity Logout.
Host Inactivity Logout	Allows you to enable the host inactivity logout and to change the timeout. See Configuring the Host Inactivity Logout.
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 IPtables. See Managing IPtables.
Login Banner Notice	Allows you to configure a login banner. See Configuring the Login Banner.
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.

Preference	Description
Remote Support	
Support Configuration	Allows you to configure remote access and notifications. Remote access enables your authorized Penguin Solutions service representative to log on to the system remotely for troubleshooting. When enabled, the everRun system can send notifications to your authorized Penguin Solutions 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 organization requires a proxy server to access the Internet and you have a service agreement with Stratus or another authorized everRun service representative. The everRun software uses proxy server information for support notification messaging and remote support access features. See Configuring Internet Proxy Settings.

Related Topics

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 **System Alias, Owner Name**, **Phone Number**, **Email**, and **Site Address**

fields.

4. Click **Save**.

Related Topics

Preferences Page

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, which updates it, if necessary.
- 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.

To check the status of a license

If the system has Internet access, use the following procedure. This procedure also automatically updates the license, if necessary. 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. For **Online License Check**, click **Check License Now**.

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

STATUS	License is activated and does not expire.
TYPE OF LICENSE	Enterprise Edition (volume)
EXPIRATION	*day, month dd*, 20*yy*, *time*
LAST CHECK	*day, month dd*, 20*yy*, *time*
ASSET ID	*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 **Offline License Check or Upload New License Key** bar to display its options, if they are not already displayed.
4. For **Upload New License Key**, click **Choose File** and navigate to the location where you saved the file. Then, click **Upload**.

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 everRun 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 **Offline License Check or Upload New License Key** bar to display its options, if they are not already displayed.
6. Under **Offline License Check via URL File**, click **Download URL File** and save the file to your mobile device or USB flash drive. If using a USB flash drive, remove it. Go to a location with Internet access.

In a location with Internet access

1. If using a USB flash drive, insert it into a USB port of the computer with Internet access.
2. Navigate to the file you saved, and click the file name.
3. A web browser opens and the Stratus license server checks the status of the license file. If necessary, a new license .key file is automatically downloaded. If using a USB flash drive, copy the new license .key file to it, and then remove the USB flash drive.
4. Return to the location with access to the console.

On a computer or mobile device with access to the 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 **Offline License Check or Upload New License** bar to display its options, if they are not already displayed.
5. For **Install an Activated License Key to the System**, 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 Penguin Solutions 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 Penguin Solutions service representative can change the limit.

9.0: ALAS_UNKNOWN

Unknown error.

Related Topics

everRun Availability Console

Preferences Page

Using the everRun Availability Console

Managing Software Updates

You can manage software updates by checking the current version number of the system software and by checking if a software update is available. You can also, optionally, enable the following:

- A message to be sent to the **Alert History** page when a system software update is available.
- An email alert (e-Alert) to be sent to a system administrator when a system software update is available.
- The system to download (though not install) the update automatically.

If you configure the system to automatically check for updates, the system checks once per day, around midnight local time. When an update is available, the system downloads it to a staging area on the system, shortly after checking for the updated software. If the download to the staging area succeeds and if configured to do so, the system sends a message to the **Alert History** page and/or an e-alert stating that the software is ready for installation. If the download fails, the update is removed.



Prerequisite: If you want system administrators to receive an e-Alert when an update is available, you must configure the mail server and e-Alerts, if these are not already configured. See Configuring the Mail Server and Configuring e-Alerts.

To manage software updates

1. Click **Preferences** in the left-hand navigation panel, to open the **Preferences** page.
2. On the **Preferences** page, click **Software Updates** (under **System**).
3. **Available System Software Updates** appears with the following information:
   - The version number of the current system software.
   - The version number of a new version of system software, if available.

   If a new version of the system software is available, click one or both of the following links, as appropriate for your needs:
   - **Download Software**–Click this link to download the available version.
   - **View Release Notes**–Click this link to view the Release Notes as well as the entire user guide for the available version.
4. **Manage System Software Updates** appears with the following options:
   - **Alert me when a System Software Update is available**–Select this option if you want a message that an update is available sent to the **Alert History** page. If you want an email sent to system administrators, informing them when an update is available, you must configure e-Alerts.

- **Automatically download System Software Updates when they become available. (Downloaded to the system only, NOT installed)**—Select this option if you want the system to download a new system software update automatically when it is available. After the software is downloaded, it is available as an upgrade kit on the **Upgrade Kits** page and you can install the software. For additional information, see the Upgrade Kits Page and Upgrading everRun Software Using an Upgrade Kit.

5. Click **Save**.

Related Topics

Alerts History Page

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 everRun system.

Warnings:

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

Notes:

1. The procedure you use to configure IP settings depends on whether the everRun system stays on the same subnet or moves to a new subnet. For instructions on how to move the system to a different subnet, access the Knowledge Base to search for the article *Moving an everRun System to a Different Subnet* (KB0013458). 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

Preferences Page

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:

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

To configure quorum servers

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

- **Port** (the default value is 4557)—Type the port number if it is different from the default.

Click **Save** to save the values.

5. Repeat steps 4 and 5 to configure a second, alternate quorum server (recommended).
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

everRun Availability Console

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

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

everRun Availability Console

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

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

Related Topics

everRun Availability Console

Preferences Page

Using the everRun Availability Console

Configuring the Mail Server

Configure the mail server to enable the everRun system to send email when, for example, someone needs to reset a password.

To configure the mail server

> **Note**: If you change any Mail Server settings, you *must* re-enter the mail-server password if authentication is enabled.

1. Click **Preferences** in the left-hand navigation panel, to open the **Preference** page.
2. Under **System**, click **Mail Server**.
3. Click the **Enable Mail Server** box. Boxes for specifying or selecting the following settings appear:
   - **SMTP Server** (required)–Enter the name of the Simple Mail Transfer Protocol (SMTP) server that your company uses to send email.
   - **Port Number** (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* (KB0012595). See Accessing Knowledge Base Articles.)
   - **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:

     > **Note**: For increased security in everRun 7.9.1.0 or higher, only the **TLS** protocol (TLS 1.2) is supported. If your mail server does not support TLS 1.2, then no outgoing emails will be sent.

     - **None** for no encryption. By default, port number 25 is used.
     - **TLS** for the Transport Layer Security (TLS) protocol. For TLS, specify 587 for **Port Number** (25 is used by default).

- **SSL** for the Secure Sockets Layer (SSL) protocol. For SSL, specify 465 for **Port Number** (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.

> **Note**: If authentication is enabled (because the **Enable Authentication** box is already checked or because you have just checked it) and you change any Mail Server settings, you *must* re-enter the mail-server password.

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

Related Topics

everRun Availability Console

Preferences Page

Using the everRun Availability Console

Configuring Users and Groups

Use the **Users & Groups** page to add, modify, or remove user accounts on the everRun system, or to grant access for Active Directory users. You can select a user and view the time when the user's password was last updated. An administrator can also use the page to force a selected user to change the user's password on the next login.

To open this page, click **Preferences** in the left-hand navigation panel and then on the **Preferences** page, select **Users & Groups** under **Administrative Tools**.

To manage local user accounts

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

To view the time when a user last changed the user's password, look at the **Last Password Update Time** column for a selected user. To force a user to change the user's password on the next login, an administrator selects the user and then clicks **Expire Password**.

For more information, see 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.

> **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

Configuring Active Directory

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. You can also select a user and view the time when the user's password was last updated, and an administrator can force a selected user to change the user's password on the next login. (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: click **Preferences** in the left-hand navigation panel to open the **Preferences** page, and then, under Administrative Tools, select **Users & Groups**.

To add a user account

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

To edit a user account

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

To force a user to change the user's password

1. Select the user whose password you want to expire.
2. Click **Expire Password**.
3. Click **Yes** in the Confirm dialog box.

To remove a user account

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

Notes:

1. You cannot delete the default **admin** account, although you should change its name and password by editing the account.
2. You must specify an email address for each user account, including **admin**, to enable the forgot password feature. If a user account does not include an email address, and the user clicks the **Forgot Password?** link on the console login page, the system sends an email to **user@example.com**.

Related Topics

Configuring Active Directory

Managing Domain User Accounts

Configuring Users and Groups

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 **Preferences** to open the **Preferences** page.
2. Under Administrative Tools, select **Users & Groups**.
3. 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**.
4. 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.Click **Deny Access**, then click **Finish**.

Related Topic

Configuring Active Directory

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 Configuring Users and Groups).

Prerequisite: If you are configuring Active Directory for the first time on everRun Release 8.*x* or higher, or upgrading from Release 7.*x* to 8.*x* or higher, you need to create new rootCA SSL certificate and install it into the Java keystore of the everRun system to enable the more secure Active Directory/LDAP over SSL support in Release 8.*x* or higher.

For detailed information about creating and installing a custom certificate for Active Directory, access the Knowledge Base (see Accessing Knowledge Base Articles) to search for the article *How to re-enable Active Directory after migrating from everRun Enterprise 7.9.x to 8.x* (KB0015895). For information about the `certificate_installer` script, see To install a custom certificate into the Java keystore.

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 **Users & Groups**.
3. Click the **Enable Active Directory** button in the lower pane.
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. Click **Add System to Active Directory**.
7. Type the **Username** and **Password** of an Active Directory Administrator in order to add this everRun system to the domain. When entering the username, enter only the username without specifying the domain. The system automatically appends the domain that you entered in step 4.
8. Click **Add**.
9. 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 **Users & Groups**.

3. Click **Remove System from Active Directory** in the lower pane.
4. Type a **Username** and **Password** that provides you with administrative privileges within the domain. When entering the username, enter only the username without specifying the domain. The system automatically appends 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 **Users & Groups**.
3. Click **Disable Active Directory** in the lower pane.

> **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

Configuring Users and Groups

Managing Domain User Accounts

Managing Local User Accounts

Preferences Page

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 () 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

Preferences Page

everRun Availability Console

Configuring Secure Connections

For security, the everRun system allows only HTTPS connections, by default. If you want to allow HTTP connections, you can configure secure connections.

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

When HTTPS connections are enabled, you can use a script to install a custom certificate into the CA bundle on the host machine. See To install a custom certificate into the CA bundle. (Alternatively, if you need to install a custom certificate into the Java keystore to enable secure connections for Active Directory or SMTP, see To install a custom certificate into the Java keystore.)

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.

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

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 you need to install a custom certificate into the **CA bundle** of the everRun system to enable HTTPS connections for the **everRun Availability Console** or the **REST API**, use the following information:

To install a custom certificate into the CA bundle

To install a custom certificate into the CA bundle of the everRun system, 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 on the host machine (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 on the host machine (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
  ```

For more information about installing a custom certificate, access the Knowledge Base to search for the article *Adding Certificates to ca-bundle.crt in everRun* (KB0013477). See Accessing Knowledge Base Articles.

The `certificate_installer` script (CA Bundle options)

Usage

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

Commands and Options for Installing Certificates Into the CA Bundle

Note: These commands and options are available only for managing certificates in the CA bundle.

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

Script Options

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

If you need to install a custom certificate into the **Java keystore** of the everRun system to enable secure connections for **Active Directory** or **SMTP**, use the following information:

To install a custom certificate into the Java keystore

To install a custom certificate into the Java keystore of the everRun system, use the `certificate_installer` script. For example, you can install a custom certificate as follows:

1. Copy a certificate to the `/tmp` folder of the host machine.
2. Issue the following command:

   ```
   certificate_installer install --java --alias ldapcert --cert
   /tmp/client.crt
   ```

For more information about creating and installing a custom certificate for Active Directory, search for the article *How to re-enable Active Directory after migrating from everRun Enterprise 7.9.x to 8.x* (KB0015895).

The `certificate_installer` script (Java keystore options)

Usage

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

Commands and Options for Installing Certificates Into the Java Keystore

> **Note**: These are the only `certificate_installer` options available for managing the Java keystore.

`install` *command_options*	Installs the custom certificate. Command options are: • `-j`, `--java`: Install a certificate into the Java keystore. • `-a`, `--alias=`*`alias`*: Specify the alias you want to assign to the certificate (recommended)

Related Topics

everRun Availability Console

Preferences Page

Using the everRun Availability Console

Configuring the Console Inactivity Logout

For security, an everRun system limits the inactivity time of a login session to the everRun Availability Console. The default timeout is 60 minutes. After 60 minutes (or other specified time) of inactivity, the everRun system automatically logs out the session. The console inactivity logout prevents a login session from remaining open indefinitely without use.

To enable the Console 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 **Console Inactivity Logout**.
3. Activate the check box next to **Enable Console Inactivity Logout**.
4. The default **Console Inactivity Logout** timeout is 60 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

Preferences Page

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

Preferences Page

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

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

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 **IPtables Security** page, you can set up, maintain, and inspect the various filter table chains and their underlying rules. You have access to the three main chains (**INPUT**, **OUTPUT**, and **FORWARD**) for applying the packet-filtering rules you need. With everRun systems, the rules are applied to 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.

You can, optionally, enable the rules to apply to the guest operating systems, in addition to the host. By default, rules apply only to the host operating system, but not to guest operating systems. When you enable rules to also apply to guests, all existing rules, imported rules, and additional newly inserted rules also apply to all guest operating systems (that is, for rules based on the same business network that has been allocated to the guest).

Notes:

1. For information on the ports that everRun software uses, see System Requirements Overview.
2. For additional information on everRun TCP and UDP ports, access the Knowledge Base to search for the article *TCP and UDP ports used by everRun 7* (KB0012595). 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 IPtables Security.

   The Enable IPtables Security window becomes gray for a few minutes. When the window is active again, Enable IPtables Security is selected

Rules are applied only to the host, by default. You can, though, apply rules to guests as well as the host.

To apply rules to guests as well as the host

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. **Apply to Host** is selected, by default:

Select **Apply to Host and Guests** to apply rules to both the host operating system and guest operating systems. The **Enable Port Management** window becomes gray for a few minutes.

When **Apply to Host and Guests** is selected, all existing rules, imported rules, and additional newly inserted rules will also apply to all guest operating systems (that is, for rules based on the same business network that has been allocated to the guest).

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 IPtables Security 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)**.
- **IP Address (starting)**–For the first IPv4 address of the range, enter an address 0.0.0.0 through 255.255.255.255 that is less than or equal to **IP Address (ending)**.
- **IP Address (ending)**–For the last IPv4 address of the range, enter an address 0.0.0.0 through 255.255.255.255 that is greater than or equal to **IP Address (starting)**.
- **IPv6 Address (starting)**–For the first IPv6 address of the range, enter an address 0000:0000:0000:0000:0000:0000:0000:0000 through ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff that is less than or equal to **IPv6 Address (ending)**.
- **IPv6 Address (ending)**–For the last IPv6 address of the range, enter an address 0000:0000:0000:0000:0000:0000:0000:0000 through ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff that is greater than or equal to **IPv6 Address (starting)**.

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

5. Newly inserted rules apply only to the host, by default. If you want the rules to apply to the host and guests, see To apply rules to guests as well as the host.
6. Click **Save** at the bottom of the page, or click **Reset** to cancel any unsaved changes, which restores rules to those of the last saved session.

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

To remove a rule

1. Click **Preferences** in the left-hand navigation panel, to open the **Preference** page.
2. On the **Preferences** page, click **IPtables Security**.

   Ensure that **Enable IPtables Security** is selected.

   (**Apply to Host** and **Apply to Host and Guests** have no effect on removing rules.)
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 IPtables Security 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 *Loading default settings....* message appears.
4. The default rules apply only to the host, by default. If you want the rules to apply to the host and guests, see To apply rules to guests as well as the host.

To import or export rules

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. Click **Import** or **Export** at the bottom of the page.
   - **Import**–The **Import/Restore IPtables Security Rules** wizard appears. Browse to and select the XML file that you want to import. All rules associated with a shared network's type within the imported XML file will be generated for each existing shared network on the system with the same type.

     After you have selected an XML file, the following message appears:

     ***Append** will reserve current rule set. Select **Overwrite** if you want to clear out all current rules.*

     Click the appropriate button:

     - **Append**–The selected XML file is appended to the existing XML file, preserving existing rules.

- **Overwrite**—The selected XML file overwrites the existing XML file, eliminating the existing rules.

- **Export**—A file explorer window appears. Browse to a location on your local system where you want to save the file of exported rules. All rules in the table are exported to an XML file that is then downloaded to the location you select.

4. Imported rules apply only to the host, by default. If you want the rules to apply to the host and guests, see To apply rules to guests as well as the host.
5. If you imported a file, click **Save** (or click **Reset** to restore the previously saved values).

Related Topics

Preferences Page

everRun Availability Console

Configuring the Login Banner

You can configure a login banner to provide customized content to the everRun Availability Console login page. For example, you can add a message.

To configure the login banner

1. Click **Preferences** in the left-hand navigation panel, to open the **Preference** page.
2. Under **Administrative Tools**, click **Login Banner Notice**.
3. Activate the **Enable Login Banner Notice** box. A box appears.

   In the box, enter the information that you want to appear on the console login page. You can, for example, type the company name or provide a message.
4. Click **Save** (or click **Reset** to restore the previously-saved values).

Related Topics

everRun Availability Console

Preferences Page

Using 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.

Prerequisite: In order for e-Alerts to function properly, you must configure the mail server. See Configuring the Mail Server.

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:
   - **e-Alerts Language**–Select a language from the pull-down menu.
   - **List of Recipients** (required)–Enter email addresses for all e-Alert recipients.
4. Click **Save** (or click **Reset** to restore the previously-saved values).

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

To generate a test alert

Click **Generate Test Alert**. The everRun software generates a test alert and sends a sample email with the subject "Test Alert" to all email recipients; SNMP sends traps to recipients of SNMP traps, if configured (see Configuring SNMP Settings); and Support Configuration sends a notification to your authorized Penguin Solutions 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

everRun Availability Console

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.penguinsolutions.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

To enable SNMP requests, perform one of the following actions:

- Enable SNMP requests from the **Preferences** page:
  - Add an SNMPv3 user who can enable SNMPv3 requests and who has read-only access to the full MIB in the everRun system.
  - Configure access control for SNMPv1 and SNMPv2 requests, where you allow no users (**Restricted**) or any user using the default public community (**Unrestricted**) to send requests.
- Customize SNMP request functionality by editing `snmpd.conf` files. You can customize access control for SNMPv1 requests and SNMPv2 requests. You can also customize the list of users for SNMPv3 requests. For information, see *To customize SNMP request functionality* (below).

To enable SNMP requests from the Preferences Page

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

4. The **List of Users for SNMP Requests (Version 3)** appears.

   If a username appears below the **List of Users for SNMP Requests (Version 3)**, the user's security level is displayed and a read-only display of the `snmpd.conf` file also appears. The user has read-only access to the full MIB. Note that the system supports only one **SNMP Requests (Version 3)** user.

   If a username does not appear, you can add an SNMPv3 user.

   To add an SNMPv3 user

   a. Click the  **Add** button, which opens the **Add a User** wizard.

   b. Enter values for the following:

      Username—The name of a user who has access to the SNMPv3 agent. The name must be unique.

      Security Level—The user's security level. Valid values are:

      - **No Authentication and No Privacy**: No security is applied to messages; messages are not authenticated or encrypted.
      - **Authentication and No Privacy**: Messages are authenticated, but not encrypted. You must enter values for **Authentication Type** and **Authentication Password**.
      - **Authentication and Privacy**: Messages are authenticated and encrypted. You must enter values for **Authentication Type**, **Authentication Password**, **Encryption Type**, and **Encryption Password**.

      When the security level includes authentication or privacy, the following fields appear:

      Authentication Type—The user's type of authentication. Valid values are:

      - **MD5**: Configure the message digest algorithm (MD5) as the user's authentication type.
      - **SHA**: Configure the secure hash algorithm (SHA) as the user's authentication type.

Authentication Password–The user's required password, which is used to generate the secret authentication key. The password must be a minimum of eight characters.

Encryption Type–The user's type of encryption. Valid values are:

- **AES**: Configure the Advanced Encryption Standard (AES) as the user's encryption type.
- **DES**: Configure the data encryption standard (DES) as the user's encryption type.

Encryption Password–The user's required password, which is used to generate the secret encryption key. The password must be a minimum of eight characters.

c. Click **Save** to save the changes.

5. Select an access option:

   Restricted (the default)–Allows no users to send SNMPv1 requests and SNMPv2 requests.

   Unrestricted–Allows any user using the default public community to send SNMPv1 requests and SNMPv2 requests.

   Customized (available when `snmpd.conf` has been manually edited by a user; see *To customize SNMP request functionality*, below)–Allows customized access.

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

To customize SNMP request functionality by editing `snmpd.conf` files

Customize SNMP request functionality by editing `snmpd.conf` files.

Customize access control for SNMPv1 requests and SNMPv2 requests by editing the `/etc/snmp/snmpd.conf` file:

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

Customize the list of users for SNMPv3 requests by editing the `/etc/snmp/snmpd.conf` and `/var/lib/net-snmp/snmpd.conf` files.

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

To enable SNMP traps

> **Notes**:
>
> 1. When you add a recipient for **SNMP Traps (Version 3)**, you need to confirm that the engine ID of the trap user on the recipient server is 0x80001370017F000001.
> 2. When you enable or modify the SNMP trap settings, generate a test alert to confirm that traps are received.

1. 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. Below the **List of Recipients of SNMP Traps (Version 3)** is a list of the trap users, and the IP address of the recipient server where the trap user exists. The everRun system sends SNMPv3 traps to the trap user on the recipient server. Add a recipient, if necessary.

   To add a recipient

   a. Click the **Add** button, which opens the **Add a Recipient** wizard.

   b. Enter values for the following:

   Recipient Address–The host name or the IPv4 address of the recipient server.

   Username–The name of a trap user on the recipient server. The name must be unique for the recipient.

Security Level—The user's security level. Valid values are:

- **No Authentication and No Privacy**: No security is applied to messages; messages are not authenticated or encrypted.
- **Authentication and No Privacy**: Messages are authenticated, but not encrypted. You must enter values for **Authentication Type** and **Authentication Password**.
- **Authentication and Privacy**: Messages are authenticated and encrypted. You must enter values for **Authentication Type**, **Authentication Password**, **Encryption Type**, and **Encryption Password**.

When the security level includes authentication or privacy, the following fields appear:

Authentication Type—The user's type of authentication. Valid values are:

- **MD5**: Configure the message digest algorithm (MD5) as the user's authentication type.
- **SHA**: Configure the secure hash algorithm (SHA) as the user's authentication type.

Authentication Password—The user's required password, which is used to generate the secret authentication key. The password must be a minimum of eight characters.

Encryption Type—The user's type of encryption. Valid values are:

- **AES**: Configure the Advanced Encryption Standard (AES) as the user's encryption type.
- **DES**: Configure the data encryption standard (DES) as the user's encryption type.

Encryption Password—The user's required password, which is used to generate the secret encryption key. The password must be a minimum of eight characters.

c. Click **Save** to save the changes.

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

7. 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)

8. 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 Penguin Solutions 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

everRun Availability Console

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 Penguin Solutions 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 Penguin Solutions 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 Penguin Solutions 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 Penguin Solutions 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 Penguin Solutions 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* (KB0012595). 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 Penguin Solutions 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

everRun Availability Console

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).
6. If the proxy server requires authentication, click the **Enable Authentication** box and type the **Username** and **Password**.

   If you do not type a password, the previous password continues to be required. If the previous password was empty and you do not enter a new password, the password remains empty.
7. Click **Save** (or click **Reset** to restore the previously-saved values).

Related Topics

everRun Availability Console

Preferences Page

Using the everRun Availability Console

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 the Audit Logs Page.)

> **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

everRun Availability Console

Using the everRun Availability Console

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

everRun Availability Console

Using the everRun Availability Console

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 Penguin Solutions 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

Related Topics

everRun Availability Console

Using the everRun Availability Console

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 Penguin Solutions service representative resolve issues with the system.



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

To create diagnostic files

1. Click **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 to indicate the amount of data and history to include in the diagnostic file:
    - Core (kernel and process core files)
    - Minimal (about a day of data, approximately 2 to 10 MB)
    - Medium (about a week of data, approximately 10 MB)
    - Large (all data except core files, approximately 50 MB)
    - Full (all data including core files, approximately 60 MB)
4. Click **Generate Diagnostic File**.
5. Upload the file to your authorized Penguin Solutions service representative, as described in Uploading a Diagnostic File to Customer Support.

Related Topics

everRun Availability Console

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 https://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

everRun Availability Console

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 Penguin Solutions 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

everRun Availability Console

Preferences Page

Using the everRun Availability Console

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, disk type, 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. The type of USB device driver is also listed.

Related Topics

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 maintenance mode.	
 Finalize	Removes a PM from the state **running (in Maintenance)**. See Maintenance Mode.
 Shutdown	Shuts down a PM. The PM transitions to **off (in Maintenance)**. See Shutting Down a Physical Machine.
 Reboot	Reboots the PM. The PM transitions to **preparing for reboot (in Maintenance)**. See Rebooting a Physical Machine.

Commands	Description
 Remove	Causes the everRun software to delete the PM from the everRun system's database, so that you can replace the PM or one of its components. See Replacing Physical Machines, Motherboards, NICs, or RAID Controllers.
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 displays the state of a failed PM as **Unreachable (Syncing/Evacuating…)**. See Recovering a Failed Physical Machine.

Related Topics

everRun Availability Console

Using the everRun Availability Console

Physical Machines Page

Physical Machine States and Activities

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

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

State	Activity	Available Commands	Description
	Powered Off	**Work On**	everRun has powered off the PM because of an excessive failure rate.
	Booting	**Finalize**	PM is booting.
	 Rebooting	**Finalize**	PM is rebooting.
	 Running	**Finalize** **Shutdown** **Reboot** **Recover** **Replace**	PM is running in Maintenance Mode. See Maintenance Mode.

Related Topics

everRun Availability Console

Using the everRun Availability Console

Physical Machines Page

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, MAC address, network, and other characteristics in the tabs of the bottom pane
- Creating, copying, exporting, importing, or restoring VMs, as described in Adding, Migrating, and Importing Virtual Machines
- Determining if your physical machines (PMs) can support VMs running Windows 11 () or if they cannot (), as described in Understanding Windows 11 Virtual Machine Differences from Other VMs
- Opening a Virtual Machine Console Session
- Changing the number of virtual CPUs (vCPUs) on a VM, changing memory allocation, storage resources, network resources, or vCPU emulation setting, as described in 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.

Related Topics

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
 Create	Launches the **Create VM** wizard. See Creating a New Virtual Machine.
 Copy	Copies an existing VM on your system to create a new VM or to create a duplicate VM for troubleshooting. See Copying a Virtual Machine.
 Import/Restore/Recover	Imports, restores, or recovers a VM, as follows: • *Imports* a VM from a set of OVF and VHD files to create a new instance of the VM. See Importing Virtualization Files. • *Restores* an existing VM on your system by overwriting the VM from a previous copy of OVF and VHD files. See Replacing/Restoring a Virtual Machine from an OVF File. • *Recovers* a VM from a set of XML and raw hard disk backup files to create a new instance of the VM. See Recovering a Virtual Machine from a Backup.
The following actions are available for use if the VM is running.	
 Mount	Mounts a USB device or a network-mounted folder (that is, a directory) to make it available to the guest operating system. You can then export a VM to the mounted location. See Mounting a USB Device or Network-mounted Folder on the everRun System.

Action	Description
 Unmount	Unmounts a mounted USB device or a network-mounted folder. See Mounting a USB Device or Network-mounted Folder on the everRun System.
 Backup	Creates a backup of a VM to a set of XML and raw hard disk backup files while the VM is running (in contrast to **Export**, which can export to OVF and VHD files only while a VM is stopped). You can optionally configure your own third-party, networked-based backup utility to collect the backup file(s) as needed. See Backing Up a Running Virtual Machine.
 Console	Opens a console for the selected VM. See Opening a Virtual Machine Console Session.
 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 actions are available if the VM is shut down or stopped.	
 Config	Launches the **Reprovision Virtual Machine** wizard. The VM must be shut down prior to launching this wizard. See Reprovisioning Virtual Machine Resources.

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

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

Related Topics

Managing the Operation of a Virtual Machine

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
	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.
	booting	Console Power Off	VM is starting. A VM remains in the **booting** state until the

State	Activity	Enabled Actions	Description
			system detects network activity from the guest operating system, at which point the VM enters the **running** state. If a VM fails to enter the **running** state, open a console window to the VM and verify that the network settings in the guest operating system are correct. If you recently imported or migrated the VM from another system, see any OS-specific procedures or troubleshooting information in Importing Virtualization Files or Migrating a Physical Machine or Virtual Machine to a System.
	 running	Console Snapshot Shutdown Power Off	VM is operating normally on redundant physical machines
	 running	Console Shutdown Power Off	VM is operating normally, but is not running on fully redundant resources.
	 stopping	Power Off Remove	VM is being shut down in response to the **Shutdown** action, or shut down because the remaining physical machine is transitioning into maintenance mode.
	 crashed		VM crashed and is restarting. If enabled, e-Alerts and support notification messages are sent.

State	Activity	Enabled Actions	Description
	✖ crashed		VM crashed too many times and exceeded its MTBF threshold. The VM is left in a crashed state until **Reset Device** is clicked. See Resetting MTBF for a Failed Virtual Machine.

Related Topics

Managing the Operation of a Virtual Machine

Virtual Machines Page

Using the everRun Availability Console

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

Related Topics

Managing Snapshots

Using the everRun Availability Console

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**
- **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

everRun Availability Console

Using the everRun Availability Console

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, disk type, 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, disk type, logical sector size, physical 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 ( ) 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

everRun Availability Console

Using the everRun Availability Console

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

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.

Related Topics

Networks Page

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

Networks Page

everRun Availability Console

Using the everRun Availability Console

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

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 check whether or not a new version of system software is available, and then download it, if available. 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.

> **Note**: You can specify that an available upgrade kit is downloaded automatically. You can also enable an email alert (e-Alert) to be sent to system administrators, notifying them when an update of system software is available. See Managing Software Updates.

To check for and download a new version of the system software

> **Note**: Your user role must be **Administrator** or **Platform Manager** to perform this procedure.

1. Click **Upgrade Kits** in the left-hand navigation panel to open the **Upgrade Kits** page.
2. Click **Check for Updates** beneath the masthead.

   A message box appears, indicating whether or not a new version of the system software is available.

3. If an update is available, the **Software Update Available** box appears, and you can click **Download Software** to download the software. You can also click **View Release Notes** to read about the update (English version).

   > **Note**: The **Upgrade Kits** page 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.

   When you click **Download Software**, the following occurs:

   - If the everRun system is connected to the Internet, a **.kit** file with the software update is downloaded directly to the system and is listed on the **Upgrade Kits** page. Various status messages appear in the **Software Update Available** box, indicating the progress of the download.
   - If the system is not connected to the Internet, the **.kit** file is downloaded to the remote management computer that is running the everRun Availability Console. Save the file to the browser's default downloads folder, or navigate to another location. You will receive an Alert (if configured) notifying you that a new version of the system software is available and that you need to upload it to the system.

4. To continue the upgrade, see Upgrading everRun Software Using an Upgrade Kit.

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

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.
2. Insert a USB medium into the primary node. On the **Physical Machines** page, check that the **USB Devices** tab lists the device.
3. In the 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.
6. If the node has more than one USB medium, you need to select a medium from the drop-down list. Then, click **Continue** (or click **Cancel** to cancel the procedure).

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

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

Related Topics

Upgrade Kits Page

4

Chapter 4: Upgrading everRun Software

To upgrade everRun software, use an upgrade kit as described in the following topics:

- To upgrade a system already running Release 8.*x* to a later version of 8.*x*, see Upgrading everRun Software Using an Upgrade Kit.
- To upgrade a system running Release 7.x to 8.x, see Upgrading from everRun Release 7.x to 8.x or Higher.

After upgrading the everRun software, optionally migrate the virtual machines to 512e storage. See Migrating Virtual Machines to 512e Storage.

Related Topics

Managing Software Updates

Upgrade Kits Page

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.

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

Prerequisites:

- To upgrade a system running Release 7.x to 8.x, see Upgrading from everRun Release 7.x to 8.x or Higher.
- For the latest information about upgrade paths and restrictions, see the everRun Release 8.1.1.0 Release Notes at http://everrundoc.stratus.com.
- 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* (KB0013295). 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.
- Ensure that any VM operations are complete, including VM backup, recover, snapshot, import, restore, and export operations.
- Shut down any unprotected VMs before upgrading the system software. If a PM contains unprotected VMs, the system automatically shuts them down before proceeding with the upgrade. For more information, see Unprotected Operation.
- 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.penguinsolutions.com/services-support/downloads/?tab=everrun. For information about how to manually install AVCLI on a remote computer see AVCLI Command Overview.

The steps are:

I. To download the upgrade kit

II. To upload the upgrade kit to the system

III. To qualify the software (optional)

IV. To upgrade the system software

I. To download the upgrade kit

When an update is available, you can download the upgrade kit that contains the new system software, if it is not already downloaded. From the **Upgrade Kits** page, click **Download Software** in the **Software Update Available** window (see the Upgrade Kits Page).

Alternatively, you can download the software from the Stratus **Downloads** page.

> **Note**: The **Upgrade Kits** page of the everRun Availability Console allows only two saved kits. If the page 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.penguinsolutions.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, click **Choose File** (in Google Chrome) or **Browse** (in Firefox or Microsoft Edge), 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 shut down, 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 **Penguin Solutions Customer Service Portal** at https://support.penguinsolutions.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.
2. Click **Yes** to continue the upgrade. The upgrade begins.

Note: After you upgrade from everRun 8.0.*x.x* to 8.1.*x.x* or higher, the everRun Availability Console updates with branding changes including a new company logo and a new color palette to reflect that Stratus Technologies is now part of Penguin Solutions. To learn more about Penguin Solutions, visit https://www.penguinsolutions.com.

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**.

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). Optionally, migrate virtual machines to 512e storage, as described in Migrating Virtual Machines to 512e Storage.

Related Topics

Managing Software Updates

Upgrade Kits Page

AVCLI Command Descriptions (see *Kit Information*)

everRun Availability Console

Using the everRun Availability Console

Upgrading from everRun Release 7.*x* to 8.*x* or Higher

Use the instructions in this topic to upgrade the system software from everRun Release 7.*x* to 8.*x* or higher. (If you need to upgrade a system already running Release 8.*x* to a later version of 8.*x*, go directly to Upgrading everRun Software Using an Upgrade Kit.)

Caution: Consider contacting your authorized Penguin Solutions service representative for assistance before upgrading your system from everRun Release 7.*x* to 8.*x*. They can help you to evaluate your system, ensure that it meets the prerequisites, locate any potential restrictions, and, if needed, discuss a strategy for the upgrade. If you encounter any problems during the upgrade process, contact your authorized Penguin Solutions service representative for assistance as soon as possible to avoid complications and potential data loss.

Prerequisites:

Before upgrading from everRun Release 7.*x* to 8.*x* or higher:

- For the latest information about upgrade paths, restrictions, and special procedures, see the everRun Release 8.1.1.0 Release Notes at http://everrundoc.stratus.com and Knowledge Base article KB0015903 in the **Penguin Solutions Customer Service Portal**. In summary, you can upgrade any everRun Release 7.9.*x.x* system directly to 8.*x* or higher. If you are upgrading from a version earlier than Release 7.9.*x.x*, you need to upgrade to Release 7.9.*x.x* first.
- Ensure that each physical node of the everRun system has at least 8 GB of RAM.
- Ensure that all PMs and VMs are in good health. 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.
- Back up the VMs and make note of any system preferences, customized configurations (for example, certificates, `snmpd.conf` modifications, or `iptables` rule changes), and logs. To export VMs on a everRun Release 7.*x* system to OVF and VHD files as a backup, see Exporting a Virtual Machine.

  In case of any failures during the upgrade process, you can use these backups to restore the system. Because the upgrade involves a migration from the earlier CentOS-based host operating system to the new Ubuntu-based host operating system, if the upgrade from Release 7.*x* to 8.*x* fails and there is no backup data, you cannot restore the system to its original status.
- In some cases, one or more VMs may be configured but stopped or intentionally shut down. If this situation applies to your system, it is important to start each of the VMs and allow them to complete any required synchronization before you start the upgrade procedure.
- When performing the upgrade, ensure that you log on to a local user account of the everRun Availability Console. To prevent an unexpected loss of access during the

upgrade process, do not use an Active Directory user account.

If your system is configured with Active Directory, you need to reconfigure Active Directory after the upgrade. See KB0015895 for information about creating a new rootCA SSL certificate and installing it into the Java keystore of the everRun system to enable the more secure Active Directory/LDAP over SSL support in Release 8.*x* or higher.

- To complete the upgrade, you must have access to the physical console (keyboard and monitor) or an out-of-band, lights-out management utility for each node in the everRun system. The following procedures discuss upgrade using only the physical console of the system.
- Record any of your current passwords for the system. The everRun Availability Console password remains the same, but, for example, the password for the host operating system on each node will revert to the default setting, and you need to change it after the upgrade to secure your system.
- If you use the snapshot feature as described in Managing Snapshots, and you delete any snapshots before upgrading, consider waiting for the coalescing operation of VM disks to complete before starting the upgrade. For more information, see Removing a Snapshot and KB0013465. Starting an upgrade while the coalescing operation is in progress could result in much more downtime than expected for the VM and the system.
- Ensure that you plan a dedicated maintenance period for the upgrade process. The upgrade requires at least 120 minutes of downtime. This time estimate does not include the time it takes to back up data from the node(s) undergoing upgrade. The upgrade also includes a loss of access to the everRun Availability Console. During the upgrade procedure, the everRun system will be completely unavailable for business connectivity.

The steps are:

I. To download the upgrade kit

II. To upload the upgrade kit to the system

III. To qualify the software

IV. To upgrade the system software from everRun Release 7.x to 8.x or higher

V. To complete post-upgrade tasks

I. To download the upgrade kit

Download the upgrade kit that contains the new system software.

> **Note**: You must manually download the everRun Release 8.*x* upgrade kit from the Stratus website and add (upload) the kit to the **Upgrade Kits** page of the everRun Availability Console. You cannot download the upgrade kit directly to the **Upgrade Kits** page.

1. Open the **Downloads** page at https://www.penguinsolutions.com/services-support/downloads/?tab=everrun.
2. Scroll down to the upgrade section and then click the **everRun 8.*x.x.x* Upgrade Kit** link to download the kit.
3. Specify 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 to the everRun system from the remote management computer that is running the everRun Availability Console.

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

1. In the everRun Availability Console, click **Upgrade Kits** in the left-hand navigation panel.
2. On the **Upgrade Kits** page, click the **Add a Kit** button beneath the masthead, which opens the **everRunKit Upload** wizard.
3. In the **Kit Upload** wizard, click **Browse** and then browse to select a .kit file.
4. After you have selected a .kit file, **Upload**, **Import**, or **Finish** (they perform the same function). A message such as **Uploading file (DO NOT CLOSE WIZARD)** appears while the file is uploading. The upload may require up to two minutes for a file stored locally, to ten or more minutes for a file stored over a network. If the upload fails, the wizard displays the message **Failed to upload file**.
5. After the upload is complete, the wizard closes and the **Upgrade Kits** page lists the state and version number of the upgrade kit. The **Qualify**, **Upgrade**, and **Delete** buttons also appear with the **Add a Kit** button.
6. If more than one upgrade kit is loaded, select the kit for 8.*x.x-xxx*.

III. To qualify the software

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

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

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

If the qualification fails, a pop-up window appears with messages indicating the cause of the failure. These messages may indicate unsupported releases, insufficient storage, partition problems, VMs that need to be shut down, 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 **Penguin Solutions Customer Service Portal** at https://support.penguinsolutions.com.

IV. To upgrade the system software from everRun Release 7.*x* to 8.*x* or higher

After verifying that your system meets the Prerequisites for upgrade, completing steps I-III of the upgrade process, and completing backups of your VMs and data, continue with the following procedure to upgrade the system software to everRun 8.*x* or higher.

1. Log on to the everRun Availability Console with a local user account (not a networked Active Directory account) for the upgrade.
2. Confirm that all PMs and VMs are healthy. In the everRun Availability Console:
   - The **Dashboard** page displays green check marks (✔) with no outstanding issues.
   - The **Physical Machines** page displays green check marks (✔) for each PM.
   - The **Virtual Machines** page displays green check marks (✔) for each VM.
3. On the **Preferences** page, click **IP Configuration**. Record the current network settings of the everRun system or take a screen capture of the **IP Configuration** page in case this information is necessary to access or restore the system.
4. Verify the health of each node in the system. Check the system logs for any memory, network, CPU, or storage errors (for example SMART disk errors) by doing the following:
   a. Open an `ssh` connection to the host operating system of a node in the everRun system at one of the IP addresses you recorded in a previous step.

b. Open logs including `/var/log/messages` in a text editor or use system commands to examine the log entries for any issues.

c. Repeat these steps on the second node in the system.

5. On the **Virtual Machines** page, if any VMs are currently stopped or intentionally shut down, verify that it is safe to introduce the VMs to your network, start each VM to verify that it is working, and allow each VM to complete any required synchronization before continuing the upgrade procedure. If insufficient resources (vCPUs or RAM) prevent all VMs from running at the same time, shut down a VM to accommodate the requirements of another VM until all VMs have been verified.

   Caution: Resolve any outstanding issues before continuing with the upgrade. If you need assistance with any repairs required to meet the health and redundancy requirements in the previous steps, contact your authorized Penguin Solutions service representative for assistance before proceeding any further.

6. On the **Virtual Machines** page, on the **Boot** tab for each VM:

   - Record the current **Auto Start Mode** setting (**Last**, **On**, or **Off**), so you can restore it after the upgrade.
   - Set the **Auto Start Mode** setting to **Off** to prevent the VM from starting during the upgrade.

7. On the **Virtual Machines** page, on the **CD Drives & USB Devices** tab for each VM, ensure that no media is connected to the VMs. If needed, disconnect the media in the guest operating system(s), and then click **Eject CD** or **Detach USB Device** in the everRun Availability Console until no devices are connected.

8. On the **Support Logs** page, create a diagnostic file for the system as described in Creating a Diagnostic File and download the file to a management PC.

   Note: When you upgrade from everRun Release 7.*x* to 8.*x*, the upgrade process does not preserve the system logs from the original Release 7.*x* (CentOS) system. It is important to create a diagnostic file to preserve the log history before upgrading the system.

9. On the **Virtual Machines** page, gracefully shut down all VMs that are running on the system.

10. On the **Upgrade Kits** page, begin the upgrade by clicking **Upgrade**.

    A **Confirm** window appears, asking you to confirm the upgrade to the selected upgrade kit.

> **Note**: Selecting the check box to enable pause has no effect. When upgrading from Release 7.*x* to 8.*x*, the system automatically pauses before upgrading the second node.

11. Click **Yes** to continue the upgrade. When the upgrade begins, **node0** reboots and **node1** shuts down in preparation for the upgrade.

    The everRun Availability Console loses its connection to the everRun system until the upgrade is complete on node0. To continue the upgrade, access the keyboard and monitor attached to the physical console of node0.

    > **Caution**: Do **not** power on node1 or start any VMs during the upgrade process. If VMs start or reboot, any business transactions occurring during this time might fail or might not be completed.

12. The physical console screen of node0 displays a text-based `grub` boot menu with a highlighted, default `Upgrade` option and waits for user interaction. Press **Enter** to proceed
13. Select a language option. The node boots from the Ubuntu OS.
14. In the text-based `Upgrade Wizard`, do one of the following:
    - To initiate the upgrade, press **Enter** and type `yes`. The upgrade begins on node0. You can monitor the upgrade progress on the physical console screen of node0. Continue with the next step.
    - To cancel the upgrade, press **1**, **c** or **C**. Node0 shuts down. If you cancel the upgrade at this point, the next time you power on node0 the `grub` boot menu displays the original `CentOS` configuration as the default boot option and automatically boots CentOS after 30 seconds if there is no user interaction.

    > **Caution**: After you initiate the upgrade on node0, you cannot stop or cancel the upgrade on node0.

    > **Caution**: If the node0 upgrade is unsuccessful for any reason, do **not** power on node1. Contact your authorized Penguin Solutions service representative for assistance before proceeding any further.

15. When the upgrade completes on node0, the node reboots. The `grub` menu displays the new `Ubuntu` configuration as the default boot option and automatically boots the Ubuntu OS after 30 seconds if there is no user interaction.
16. After the Ubuntu OS starts, the console displays the node0 and system IP addresses after a few minutes. If needed, record the system IP address to access the everRun Availability Console in the next step. The system IP address is the second IPv4 address.

    ```
    everRun
    IPv4 address 10.201.128.201 <-- node0 IPv4 address
    IPv4 address 10.201.128.200 <-- system IPv4 address
    IPv6 address 3d00:feed:face:1083:225:64ff:fe8d:1b6e
    IPv6 address fe80::225:64ff:fe8d:1b6e
    ```

17. If the everRun Availability Console is still running on the remote management computer, refresh the web browser to view the current status of the system; or, if needed, open a web browser, load the system IP address, and log on to the everRun Availability Console.

    > **Note**: After you upgrade to everRun 8.1.*x.x* or higher, the everRun Availability Console updates with branding changes including a new company logo and a new color palette to reflect that Stratus Technologies is now part of Penguin Solutions. To learn more about Penguin Solutions, visit https://www.penguinsolutions.com.

    > **Note**: After the system upgrades node0, but node1 has not been updated yet, the nodes are running different types of the system software. During this time, the everRun Availability Console masthead might display healthy, critical, warning, or other health states, which is expected because the nodes are unable to communicate with each other. VMs, volumes, storage groups, and networks will also have warnings (⚠) because they are currently in simplex mode.

18. In the everRun Availability Console, confirm both of the following:
    - On the **Physical Machines** page, verify that only node0 is currently **running** and displayed as **node0 (primary)**, whereas node1 is **lost**.

- In the everRun Availability Console masthead, verify that the system is now running everRun Release 8.*x*. For example:

```
ocean.abc.com
IP: 123.109.50.34 | Asset ID: ee-12345
Version: 8.x.x-xxx
```

If these conditions are met, continue with the next step; otherwise, contact your authorized Penguin Solutions service representative for assistance.

19. Start the VMs one at a time and verify that they operate normally. For each VM, do the following:

> **Caution**: It is important to verify the VMs, but note that when you start a VM, it acquires new data from the network, and this copy of the VM becomes the most up-to-date version of the VM. If you roll back to the original everRun Release 7.*x* system, you might need to re-export the VMs before rolling back to preserve any changes.

    a. Start the VM and verify that it reaches the **running** state. If the VM appears stuck in the **booting** state, verify its network connections and retry.
    b. Log on to the console of the VM.
    c. Evaluate the VM as it is admitted to its networks.
    d. If the VM is operating normally, reconfigure its **Auto Start Mode** as desired.

> **Caution**: If the node0 upgrade is unsuccessful for any reason, do **not** power on node1. Contact your authorized Penguin Solutions service representative for assistance before proceeding any further. If you proceed with the next step, the upgrade process overwrites node1, and you cannot roll back to the original, CentOS-based host operating system without fully reinstalling the everRun software.

20. After verifying that node0 has been upgradefd successfully, manually power on node1 to continue the upgrade.
21. The physical console screen of node1 displays a text-based `grub` boot menu with a highlighted, default `Upgrade` option and waits for user interaction. Press **Enter** to proceed
22. Select a language option. The node boots from the Ubuntu OS.
23. In the text-based `Upgrade Wizard`, do one of the following:

- To initiate the upgrade, press **Enter** and type **yes**. The upgrade begins on node1. You can monitor the upgrade progress on the physical console screen of node1. Continue with the next step.
- To cancel the upgrade, press **1**, **c** or **C**. Node1 shuts down. If you cancel the upgrade at this point, the next time you power on the system the `grub` boot menu displays the original `CentOS` configuration as the default boot option and automatically boots CentOS after 30 seconds if there is no user interaction.

24. When the upgrade completes on node1, the node reboots. The `grub` menu displays the new `Ubuntu` configuration as the default boot option and automatically boots the Ubuntu OS after 30 seconds if there is no user interaction.
25. After the upgrade completes on node1, it takes a few minutes for node1 to join the system.
    - On the **Physical Machines** page, node0 will be **running** with a warning () and node1 will be **(in Maintenance)** mode.
    - After node1 joins the system, both nodes will be in the **running** state. The warnings () might continue until the system synchronizes all data been the nodes.
    - After all volumes have synchronized, all warnings disappear, and all components should have green check marks ().

    Confirm that there are no critical alerts in the masthead or on the **Dashboard** page of the everRun Availability Console. Resolve any remaining issues. When the system is healthy, the upgrade is complete.
26. Ensure that the system is running correctly by checking all of the features.

V. To complete post-upgrade tasks

1. Remove the obsolete VirtIO ISO file.

   After the upgrade, the everRun Availability Console might still contain an old VirtIO driver ISO file from everRun Release 7.*x*. You must remove the ISO file as follows:

   a. In the everRun Availability Console, open the **Virtual CDs** page.

   b. On the **Virtual CDs** page, remove the **virtio-win-0.1.171_rev1** ISO by selecting the ISO and clicking **Remove**.

> **Caution**: Do **not** remove the **virtio-win-0.1.190_rev2** or higher ISO associated with everRun Release 8.*x* or higher.

2. Update the VirtIO drivers on all Windows-based VMs, as described in Updating the VirtIO Drivers (Windows-based VMs). Ensure that you install the VirtIO drivers from the new **virtio-win-0.1.190_rev2** or higher ISO associated with everRun Release 8.*x* or higher.
3. Reset the `root` password on each node in the system.

   After the upgrade, the password on each node resets to the default `KeepRunning`. Log on to the host operating system of each node and reset the `root` password to secure your system. If the node does not prompt for a new password upon login, issue the `passwd` command to change the password.
4. If applicable, restore your Active Directory configuration. Create a new security certificate, install the certificate, and reconfigure Active Directory to allow secure logins as described in KB0015895.
5. Restore any other security certificates, as needed.
6. Restore any host operating system customizations.
7. Continue to configure and test the everRun system, as needed.
8. Optionally, migrate virtual machines to 512e storage, as described in Migrating Virtual Machines to 512e Storage.

Migrating Virtual Machines to 512e Storage

everRun Release 7.8.0.0 or higher includes support for presenting volumes to the guest operating systems as the 512e disk type. If there are 512e disk drives in both physical machines (PMs) of an everRun system, and the applications running in the virtual machines (VMs) could benefit from 512e drive performance, it might be possible to reconfigure the system to take advantage of the 512e drives.

When you upgrade to everRun Release 7.8.0.0 or higher, existing storage groups are not automatically reconfigured for 512e support, and they cannot be reconfigured in place; however, it might be possible to assign new or unused 512e drives to a new 512e storage group to make a pool of 512e storage available for the VMs to use, as follows:

- If your system has empty disk slots, insert additional 512e disk drives, create storage groups for data with the 512e disk type, and allocate the storage as described in this topic under **To migrate compatible systems to 512e storage**.

- If your system has no empty disk slots for additional 512e drives, run the 512e compatibility script to determine if the system has unused 512e disk drives that you can assign to new 512e storage groups, or if you would need to rebuild the everRun system to free up 512e disk drives. You can run the script as described in this topic under **To determine system compatibility for migration to 512e storage**.

If you cannot add new 512e disk drives, and the system has no unused 512e drives (for example, if it has 512e drives that are currently in use by 512n storage groups), it might be possible to rebuild the everRun system to free up 512e drives, but the rebuild procedure is time consuming and involves significant downtime. If applicable, see KB0014930.

To determine system compatibility for migration to 512e storage

After upgrading an everRun system to Release 7.8.0.0 or higher, you can run a compatibility script to determine if the system has any unused 512e drives or drives that can be reconfigured with a system rebuild:

1. Before running the script, activate all new or foreign drives in both PMs, as described in Activating a New Logical Disk.
2. Log on to the host operating system of each PM in the everRun system as the `root` user. (For details, see Accessing the Host Operating System.)
3. Run the compatibility script in the host operating system of each PM:

   ```
   # /opt/ft/sbin/find_available_512e_drives
   ```

4. Examine and compare the output.

For example, if migration is possible now with unused 512e drives, the output is similar to the following:

```
# /opt/ft/sbin/find_available_512e_drives
512e drives in this system:
/dev/sdb ('Logical Disk - 1')

512e drives that can be moved to a 512e storage group now:
/dev/sdb ('Logical Disk - 1')
```

If migration is possible only with a system rebuild, the output is similar to the following:

```
# /opt/ft/sbin/find_available_512e_drives
512e drives in this system:
/dev/sdb ('Logical Disk - 1')
```

```
No 512e drives can be moved to a 512e storage group now

512e drives that can be moved to a 512e storage group with a
system rebuild:
/dev/sdb ('Logical Disk - 1')
```

Compare the output from both PMs and determine the next steps:

- If both PMs have unused 512e disk drives that you can use without rebuilding the system, then create a new 512e storage group, assign the drives to the storage group, and copy or move the VMs to the storage group as described in this topic under **To migrate compatible systems to 512e storage**.
- If the only way to make 512e disk drives available is to rebuild the everRun system, then refer to KB0014930.

To migrate compatible systems to 512e storage

If your everRun system contains new or unused 512e disk drives, and the compatibility script confirms that you can move these drives to a 512e storage group now without a system rebuild, use the following steps to migrate your VMs to 512e storage:

> **Note**: While migrating to 512e storage, observe the planning information and restrictions described in Storage Requirements. 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.

1. Back up the VMs.
2. For Windows-based guest operating systems only, if you have not already done so, download and update the VirtIO drivers as described in Updating the VirtIO Drivers (Windows-based VMs). At a minimum, you must update the storage drivers to support 512e volumes.
3. Create new, empty storage group(s) with the 512e disk type, as described in Creating a New Storage Group.
4. Assign the new logical disk(s) (paired physical disk drives) to the storage group(s) that you created, as described in Assigning a Logical Disk to a Storage Group.
5. Create new VMs in the new 512e storage group(s), as described in Creating a New Virtual Machine, or move existing VMs from the current storage groups by doing one of the following:

- Copy existing VMs to the new 512e storage group(s), as described in Copying a Virtual Machine. (If your system lacks the disk space to copy the VMs, try the next option.)
- Export existing VMs as described in Exporting a Virtual Machine, and then import them to the new 512e storage group(s) as described in Importing Virtualization Files.

5

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. For information, see Responding to a Failed Logical Disk

Related Topics

Logical Disks and Physical Disks

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.

Related Topics

Logical Disks and Physical Disks

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 everRun software is able to access logical disks that the RAID controller presents to the operating system. When the everRun software recognizes a new logical disk, it performs one of the following actions:

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

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

Prerequisite: Note any relevant requirements in Storage Requirements.

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

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.

Notes:

- 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.
- Select a disk type (512e, 512n, or 4k) that is most optimized to the logical disks in the system and the VMs that you want to run. You cannot change this disk type after you create the storage group, and you can assign logical disks to the storage group only if they are compatible with this disk type. See Storage Requirements for more information.

To create a new storage group

1. On the **STORAGE GROUPS** page, click the **Create** button. The **New Storage Group** dialog box appears.
2. In the **Name** box, type the name of the new storage group.
3. For **Disk Type**, select the disk type for the storage group: **512e**, **512n**, or **4k**.
4. 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.

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.

Notes:

- The logical disk that you want to assign must be empty. Its **Size Used** value must be zero.
- If you want to reassign a logical disk that is currently in use, you must delete all volumes from that logical disk before you can move it to a new storage group.
- Note any relevant requirements in Storage Requirements.

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. (The drop down box lists only the storage groups that are compatible with the disk type of the logical disk.)
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

6

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 (**node*x* (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.
2. 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.

Note: Shut down any unprotected VMs before entering a PM into maintenance mode. If the PM contains unprotected VMs, the system automatically shuts them down before initiating maintenance mode. For more information, see Unprotected Operation.

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

everRun Availability Console

Managing Physical Machines

Physical Machines and Virtual Machines

Physical Machines Page

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.
3. Select the appropriate PM (node0 or node1) and then click **Work On**, which changes the PM's **Overall State** to **Maintenance Mode** and the **Activity** state to **running (in Maintenance)**.
4. Click **Reboot**. As the PM reboots, the **Activity** state displays:
   - **preparing for reboot (in Maintenance)**
   - **rebooting (in Maintenance)**
   - **booting (in Maintenance)**
   - **running (in Maintenance)**.
5. To exit the PM from maintenance mode and make it available for running virtual machines, click **Finalize**.

Related Topics

Maintenance Mode

everRun Availability Console

Managing Physical Machines

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 shutdown a PM, as described in the procedure below.
3. The everRun system is not fault tolerant when you shut down a PM. For continuous uptime, bring an offline PM back into service as soon as possible.

To shut down a PM

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 ✖**off (in Maintenance)**. You must manually restart the PM.

Related Topics

Maintenance Mode

everRun Availability Console

Managing Physical Machines

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.

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 of an 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 re-creates 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 Penguin Solutions 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.
2. Select the appropriate PM (node0 or node1) and then click **Work On**, which changes the PM's **Overall State** to **Maintenance Mode** and the **Activity** state to **running (in Maintenance)**.
3. After the PM displays **running (in Maintenance)**, click **Recover**.
4. When prompted to select the type of repair, click **PXE PM Recover - Preserve Data**.

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

5. Click **Continue** to begin the recovery process. The system reboots the target PM in preparation for the system software reinstallation.
6. 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.
7. When the software installation is complete, the target PM reboots from the newly installed system software.
8. As the target PM boots, you can view its activity on the **Physical Machines** page of the everRun Availability Console. The **Activity** column displays the PM as **(in Maintenance)** after the recovery is complete.
9. If applicable, manually reinstall applications and any other host-level software, and reconfigure the PM to match your original settings.

10. When you are ready to bring the target PM online, click **Finalize** to exit maintenance mode. Verify that both PMs return to the **running** state and that the PMs finish synchronizing.

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

To recover a PM (with DVD/USB installation)

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.
2. Select the appropriate PM (node0 or node1) and then click **Work On**, which changes the PM's **Overall State** to **Maintenance Mode** and the **Activity** state to **running (in Maintenance)**.
3. After the PM displays **running (in Maintenance)**, click **Recover**.
4. When prompted to select the type of repair, click **DVD/USB PM Recover - Preserve Data**.

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

5. Click **Continue** to begin the recovery process. The system shuts down the target PM in preparation for the system software reinstallation.
6. Insert the bootable media or mount the ISO image on the target PM, and then manually power on the PM.
7. 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**:
>
> 1. If you are not sure of which port to use, use the arrow keys to select one of the ports, and click the **Identify** button. The LED on the selected port will then flash for 30 seconds, allowing you to identify it. Since the LED may also flash due to activity on that network, you can 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.

> **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.

14. 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 everRun 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.
19. 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

everRun Availability Console

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 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)
- Adding, Migrating, and Importing 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 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 has 4 cores and 2 threads per core (8 vCPUs) and a second PM (in that system) has 8 cores and 2 threads per core (16 vCPUs), the total number of vCPUs available is 8 vCPUs (fewest threads of either PM).
- The number of vCPUs available for all VMs is equal to the total number of vCPUs available on the everRun system minus the number of vCPUs allocated to the everRun system software. (You 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.
- The everRun software allows the over-provisioning of vCPUs. If the number of free vCPUs on the **System** page is less than zero, you have over-provisioned the vCPUs; the console indicates this and displays an estimate of the degree to which vCPUs have been over-provisioned.
- The over-provisioning of vCPUs does not prevent you from creating or starting VMs; however, it is best to avoid running the system in an over-provisioned state.

Considerations When Over-Provisioning Virtual CPUs

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

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

- The peak vCPU resources consumed by the combined VMs does not exceed the physical resources of the everRun system.
- One or more VMs are used at different times (such as off-peak backups).
- One or more of the VMs will be stopped while the other is running, for example, during VM upgrades or VM point-in-time backup or recovery.
- Peak total CPU use by VMs will not affect service level agreements or required response times.
- Each VM's CPU use is well understood, and its application(s) are not prone to resource leaks. When CPUs are over-provisioned, a leak in one VM can affect the performance of other VMs.

Related Topics

System Requirements Overview

Adding, Migrating, and Importing 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

Adding, Migrating, and Importing 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

  For best performance, use disks with the 4K sector size. When you create each storage group, ensure that you specify the disk type that is most optimized to the logical disks in the system and the VMs that you want to run. This allows you to maximize the performance of 512e and 4K disks, when they are compatible with the VMs that you want to create or import:

  - A storage group with the 512n or 512e disk type provides the sector size of 512B for its VM volumes.
  - A storage group with the 4K disk type provides a sector size of either 4K or 512B, selectable for each of its VM volumes.
  - If a storage group with the 512e or 4K disk type provides a volume with the 512B sector size, it is presented to the VM as a volume with the 512e disk type.

  Note that the boot volume for each VM must be 512B, regardless of the storage group's disk type. 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

Adding, Migrating, and Importing 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

Adding, Migrating, and Importing Virtual Machines

Managing Virtual Machine Resources

Adding, Migrating, and Importing Virtual Machines

You can add new virtual machines (VMs) on a everRun system by:

- Creating a new VM without an existing source VM or physical machine (PM). See Creating a New Virtual Machine.
- Migrating an existing VM or PM directly over the network.
- Importing an Open Virtualization Format (OVF) file or Open Virtualization Appliance (OVA) file from an existing VM.

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

To back up a running VM for the purpose of recovering the VM in the event of a failure, see Backing Up a Running Virtual Machine and Recovering a Virtual Machine from a Backup.

To migrate or import a VM from another system, from a VCD, or to restore a VM on the same everRun 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 Virtualization Files

  Use the everRun Availability Console to import OVF and VHD files from another everRun system to the everRun system, or to import OVF and VHD files or an OVA file from a VMware vSphere-based system to the everRun system.

- Importing Virtual Machines from a VCD

  Pack virtualization files into an ISO, upload the ISO to the system as a VCD using the everRun Availability Console, and import the VM using the Representational State Transfer (REST) API. This option is useful when bandwidth issues prevent remote VM import.

- 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.

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.

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.

When you create a new VM, the everRun system assigns the Skylake CPU model by default. The default emulation persists across VM shutdown, restart, and upgrade, unless you explicitly override it. For information about available emulations and instructions for overriding the default CPU emulation setting, see Reprovisioning Virtual Machine Resources.

Prerequisites:

- Review the prerequisites and considerations for allocating CPUs, memory, storage, and network resources to the VM, as listed in Planning Virtual Machine Resources as well as Virtual Machine Recommendations and Limits.
- You can create VMs that run supported guest operating systems and boot interfaces. These are listed in Tested Guest Operating Systems.
- If you create a VM with the Windows 11 guest operating system, you must select the UEFI boot interface to support its strict security requirements. For more information, see Understanding Windows 11 Virtual Machine Differences from Other VMs.
- You can select a remote or local ISO or a bootable virtual CD (VCD) as the source that the VM boots from. For a remote ISO, you must have a URL or path name for the repository; and for a remote ISO on a shared network drive, you must have a user name and password. If you need a bootable VCD of the Windows or Linux installation media, create it as described in Creating a Virtual CD. The bootable VCD must be a single CD or DVD. Multiple CDs or DVDs are not supported.
- Ensure that both PMs of the everRun system are online; otherwise, the system cannot properly create the VM.

To create a new VM

1. On the **Virtual Machines** page (see Virtual Machines Page), click **Create** to open the **VM Creation** wizard.
2. On the **Name, Description, and OS** page:
   a. Type the **Name** and an optional **Description** for the VM as you want them to appear in the everRun Availability Console.

      The VM name must meet the following requirements:

      - Must start with a word or a number.
      - Cannot be blank.
      - Cannot include the following special characters: spaces, & (ampersand), < (left angle bracket), > (right angle bracket).

- Cannot use hyphenated prefixes such as Zombie- or migrating-.
- Has a maximum of 64 characters.

b. For **Boot Interface**, select one of the following:

- **BIOS**–Basic Input/Output System
- **UEFI**–Unified Extensible Firmware Interface

> **Notes**:
>
> 1. Ensure that the guest operating system supports the **Boot Interface** that you select; otherwise, the guest operating system cannot boot properly. For a list of guest operating systems and boot interfaces that are supported on everRun systems, see Tested Guest Operating Systems.
> 2. You can set the **Boot Interface** only when creating a VM. You cannot modify the setting later.

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.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 **Protection Level** page, 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 PMs. 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.
- **Unprotected (NONE)**—Provides no redundancy or failover. The VM runs on only one PM; if that PM goes offline for any reason, the VM stops running until the PM is back online. Use unprotected operation only if the VM has specific requirements, for example, if software licensing restrictions prevent your applications from running on both PMs for redundancy.

For more information about these levels of protection, see Modes of Operation.

After selecting the level of protection, click **Next**.

4. On the **vCPUs and Memory** page:
   a. Specify the number of **vCPUs** and the amount of **Memory** to assign to the VM. For more information, see Planning Virtual Machine vCPUs and Planning Virtual Machine Memory.
   b. Click **Next**.
5. On the **Volumes** page:
   a. Type the **Name** of the boot volume as it will appear in the everRun Availability Console.

      The volume name must meet the following requirements:

      - Must include only the valid characters a-z, A-Z, 0-9 , +, _, and . (dot).
      - Cannot be only . or .. (dots).
      - Cannot be snapshot or pvmove.
      - Cannot contain any of the following reserved strings: _cdata, _cmeta, _corig, _mlog, _mimage, _pmspare, _rimage, _rmeta, _tdata, _tmeta, or _vorigin.

   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 andPlanning 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 **Volume 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**.

6. On the **Networks** page, select the shared networks to attach to the VM (for more information, see Planning Virtual Machine Networks). You can also enable (or disable) the network and specify the MAC address. To continue, click **Next**.

7. On the **Creation Summary** page:

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

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

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

   The **VM Creation** wizard displays progress of the creation and opens the console window, if applicable. When the console window opens, it may take up to a minute for the console to connect to the VM.

8. For Windows-based VMs, when the VM console opens, click inside the console window and be prepared to press any key to run **Windows Setup** from the VCD or remote ISO.

   ```
   Press any key to boot from CD or DVD...
   ```

   For Windows-based VMs with the UEFI boot type, you need to press a key within one or two seconds; otherwise, the **UEFI Interactive Shell** appears. If this happens, you can recover and run **Windows Setup** as follows:

   a. In the **UEFI Interactive Shell**, at the `Shell>` prompt, type `exit` and press **Enter**.

      ```
      Shell> exit
      ```

b. Use the arrow keys to select **Continue**, and press **Enter**.

```
Select Language

Device Manager
Boot Manager
Boot Maintenance Manager

Continue
Reset
```

c. As the VM restarts, press any key to run **Windows Setup** from the VCD or remote ISO.

```
Press any key to boot from CD or DVD...
```

d. If you miss pressing any key, and the **UEFI Interactive Shell** is displayed again, repeat steps a-c.

9. 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.
10. 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

Reprovisioning Virtual Machine Resources

Adding, Migrating, and Importing Virtual Machines

Managing Virtual Machine Resources

Managing the Operation of a Virtual Machine

Copying a Virtual Machine

Copy a virtual machine (VM) if you want to clone an existing VM on your 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 Adding, Migrating, and Importing Virtual Machines.)

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

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

Notes:

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

Prerequisites:

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

To copy a VM on the everRun system

1. On the **Physical Machines** page (see the Physical Machines Page), verify that both PMs are in the **running** state and that neither PM is in maintenance mode or in the process of synchronizing.
2. On the **Virtual Machines** page (see the Virtual Machines Page), select the VM that you want to copy and click **Shutdown**.
3. When the VM has stopped, click **Copy** to open the **Copy** wizard.
4. On the **Name, Description, and Protection** page:
    a. Type the **Name** and an optional **Description** for the VM as they will appear in the everRun Availability Console.
    b. Select the level of protection to use for the VM:
        - **Fault Tolerant (FT)**
        - **High Availability (HA)**

        For information about these levels of protection, see Creating a New Virtual Machine and Modes of Operation.
    c. Click **Next**.
5. On the **vCPUs and Memory** page:
    a. Specify the number of **vCPUs** and the amount of **Memory** to assign to the VM. For more information, see Planning Virtual Machine vCPUs and Planning Virtual Machine Memory.
    b. Click **Next**.
6. On the **Volumes** page, you can:
    - Type the volume **Name**.
    - Specify the **Container Size** and **Volume Size** of each volume to make room for snapshots.
    - Assign each volume to a **Storage Group**.

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**.

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

After the copy process is complete; the everRun system may continue to synchronize data between PMs to enable 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

Adding, Migrating, and Importing Virtual Machines

Managing Virtual Machine Resources

Managing the Operation of a Virtual Machine

Backing Up a Running Virtual Machine

> **Note**: The virtual machine (VM) backup and recovery feature is a Technology Preview only and is under development for potential availability in a future release. Access to the feature requires a separate license and is disabled by default. The feature is not supported nor intended for use in production environments. Contact your authorized Penguin Solutions service representative for more information.

Back up a running VM to save a backup of the VM to a network-mounted folder or to a USB device. Backing up a VM on the everRun system makes the VM image available for recovery on the same everRun system or on another everRun system, as described in Recovering a Virtual Machine from a Backup. You can also optionally configure a third-party, network-based backup solution to collect the files from the backup location on a regular basis for additional redundancy. Alternatively, you can create a backup by exporting a VM from a system as standard Open Virtualization Format (OVF) and Virtual Hard Disk (VHD) files, as described in Exporting a Virtual Machine; however, you can export a VM only while the VM is stopped.

Prepare for backing up a VM by inserting a USB device or by creating a network-mounted folder to store backups 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 a backup in the everRun Availability Console, the everRun system saves the VM as Extensible Markup Language (XML) and raw hard disk data files.

Notes:

- By default, all backups are considered *crash-consistent* backups, similar to the system crashing or being powered off unexpectedly. Any data that is cached or in flight at the time of the backup might be lost unless you close all running applications for the duration of the backup, or shut down the guest VM and export it (see Exporting a Virtual Machine) as an alternative method for the VM image backup.
- The time required for the backup 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 restoring it, remember to set a different MAC address and IP address for the VM when you recover it on the target system.
- If the everRun system switches from the primary PM to the secondary PM during a backup, the process fails. This does not affect the continuous uptime of the system. You can delete the partially backed up files from the network-mounted folder and back up the VM again.
- The maximum size of VM volume data that you can back up to an external file share or USB device formatted as a FAT or VFAT file system is 4 GB, even if the total size of the target device is much larger. If you attempt to back up data more than 4 GB to a FAT or VFAT file system, the backup will fail.
- You cannot back up an unprotected VM. For more information, see Unprotected Operation.

Prerequisites:

- Ensure that the VM is in the **running** state and healthy. You cannot back up a VM if it is **stopped**.
- Prepare the backup destination:
  - If you are using a USB device, insert it into the current primary node of the system (displayed as **node*n* (primary)** on the **Physical Machines** page). Confirm that 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.
  - 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 VM backup. 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 back up the VM.
  - Ensure that the backup destination has enough storage space for the VMs that you want to back up.

To back up a VM

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 the VM that you want to back up and ensure that it is in the **running** state and healthy. See the Virtual Machines Page.
4. With the VM selected, click **Backup** in the bottom pane.
5. In the **Backup** dialog box, 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 **Backup** dialog box displays the mounted device URL in green. To change it, click the **Change** button.

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

  The backup 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\MyBackupVMs**).

- **Mount device via NFS**

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

- **Mount USB**

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

6. For **Backup Path: `/mnt/ft-export`:**, type the path of the location where you want the VM backup files to be stored. For example, if you want to back up the VM to a new folder named `ocean1`, type `ocean1`.
7. Click **Mount**.

   If the mount succeeds, the repository appears under **Device URL** and the **Backup VM** button becomes active; otherwise, an alert appears.
8. Click **Backup VM** to back up the VM.

You can monitor the **Backup Status** in the **Summary** tab for the VM that you are backing up. When the process is complete, the status changes to **Backup completed successfully**.

If you need to cancel the backup, click **Cancel** next to the **Backup progress** percentage. A dialog box opens, asking you to confirm the cancellation. Click **Yes** to cancel.

The everRun system backs up the volumes first, then it backs up the XML metadata file. You can confirm that the process is finished when the XML file appears in the folder.

After the backup process, if you need to recover the VM, see Recovering a Virtual Machine from a Backup.

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 backup process.

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

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

Related Topics

Attaching a USB Device to a Virtual Machine

Adding, Migrating, and Importing Virtual Machines

Managing Virtual Machine Resources

Managing the Operation of a Virtual Machine

Recovering a Virtual Machine from a Backup



Note: The virtual machine (VM) backup and recovery feature is a Technology Preview only and is under development for potential availability in a future release. Access to the feature requires a separate license and is disabled by default. The feature is not supported nor intended for use in production environments. Contact your authorized Penguin Solutions service representative for more information.

Recover a VM to recover a backup of the VM from a network-mounted folder or from a USB device. Recovering a VM from a backup that was created on the everRun system as described in Backing Up a Running Virtual Machine creates a new, duplicate VM on the same everRun system or on another system. (Alternatively, you can import or restore a VM to a system from standard Open Virtualization Format (OVF) and Virtual Hard Disk (VHD) files, as described in Importing Virtualization Files and Replacing/Restoring a Virtual Machine from an OVF File.)

Recovering a VM creates an identical VM with the same SMBIOS UUID, system serial number, and MAC addresses, if provided in the VM backup, that your guest operating system and applications may require for software licensing. The hardware ID, though, of the recovered VM is unique.

You can recover a VM only if you have previously backed up the VM (as described in Backing Up a Running Virtual Machine) from an everRun system to files on a supported network share or a USB device. 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 recover the Extensible Markup Language (XML) and raw hard disk data files.

Notes:

- You can recover a VM only from backup files created with the **Backup** feature on an everRun system. You cannot recover a VM from backup files created on a third-party system. If needed you can import OVF and VHD files that have been exported from a third-party system.
- When recovering a VM, the system attempts to preserve the hardware ID and MAC addresses of all network interfaces.
- Recover a VM only if you are specifically trying to recover a particular instance of an everRun VM and the recovered VM will be the only copy of this VM running across all everRun servers in your network.
- The time required to recover 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 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 recover the VM so that you can reconfigure the containers of the new VM after the recovery 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 recovering a VM, the recovery 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 recover them again.

Prerequisites:

- Before you recover a VM from an everRun system, use the everRun Availability Console on the source everRun system to back up the VM to a supported network share or a USB device. 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 recover the VM.
- Both PMs of the everRun system must be online for the recover process to function properly.

To recover a VM

1. Log on to the everRun Availability Console on the target everRun system.
2. On the **Physical Machines** page (see the Physical Machines Page) of a system configured with two nodes, verify that both PMs are in the **running** state and that neither PM is in maintenance mode or in the process of synchronizing.
3. Mount the USB device or network share on the everRun system as described in Mounting a USB Device or Network-mounted Folder on the everRun System.
4. On the **Virtual Machines** page (see Virtual Machines Page), select the VM that you want to recover in the upper panel.
5. Click **Import/Restore/Recover** to open the **Import/Restore/Recover Virtual Machine** wizard.
6. Select **Recover**, and then click **Next**.
7. Select one of the following:
   - **Recover from USB**–Recovers the VM from a USB device mounted on the everRun system.

     Click **Next** and then select a partition from the pull-down menu. Click **List backup files** and select the appropriate file from the pull-down menu. You can optionally search for a file by entering the file name or partial file name in the *Search Files* box. The box lists backup files that have names matching the name entered in the box, and that reside in various directories:
     - With the parent (root) directory as the search directory, the listed files reside in sub-directories in addition to the parent (root) directory.

- With a sub-directory as the search directory, the listed files reside in the parent (root) directory in addition to the sub-directory.

- **Recover from remote/network Windows Share(CIFS/SMB)**–Recovers the VM from a Windows share on your local network.

  Click **Next** and 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 backup files** and select the appropriate backup file from the list.

- **Recover from remote/network NFS**–Recovers the VM from an NFS share on your local network.

  Click **Next** and for **Repository**, enter the URL of the remote system in the format ***nnn.nnn.nnn.nnn/folder_name*** (do not include **http://** or **https://**).

  Click **List backup files** to display a list of all files in the remote folder. Select the appropriate backup file. You can optionally search for a file by entering the file name or partial file name in the *Search Files* box, or you can reorganize the list by clicking a column heading (*Name*, *Date Modified*, or *Size*). Click the file name to select the file.

When you finish selecting a backup file, click **Next** to continue.

8. Select the **Protection Level** to use for the VM:

   - **Fault Tolerant (FT)**
   - **High Availability (HA)**
   - **Unprotected (NONE)**

   For information about these levels of protection, see Creating a New Virtual Machine and Modes of Operation.

9. Click **Recover** to begin recovering the VM.

   You can monitor the **Recovery Status** in the **Summary** tab for the VM that you are recovering. When the process is complete, the status changes to **Recovery completed successfully**.

   > **Note**: Recovered volumes begin to appear on the **Volumes** page of the everRun Availability Console while the recovery process is still in progress. Do not attach or remove any of these recovered volumes until the recovery is complete; otherwise, the recovery process fails.

10. 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 recovered VM is functioning properly, the recovery 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 recovered VM and its associated volumes may be marked with warning symbols until the data has been synchronized and the VirtIO drivers are running.

Troubleshooting

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

To clean up after a canceled or failed recovery process

In the everRun Availability Console on the target system, remove the recovered VM and any volumes associated with the recovered VM, if present.

Related Topics

Adding, Migrating, and Importing 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) or Open Virtualization Appliance (OVA) file to a system, as summarized in Adding, Migrating, and Importing Virtual Machines.)

The procedures below describe how to migrate a PM or VM over the network: download the *P2V client* (**virt-p2v**) ISO file, boot the P2V client ISO file on the source PM or VM, and then use the client to configure, initiate, and monitor the secure network transfer from source side. No configuration steps are required on the system until after the migration is complete, but you can confirm that the migration is in progress on the

Volumes page of the 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). For additional information on backing up VMs or PMs, see Security Hardening.

Notes:

- You can migrate VMs only if they are running supported guest operating systems, as described in Tested Guest Operating Systems.
- 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.
- While migrating a VM from an everRun or ztC Edge system, it is normal if the source system displays the alert "The VM *name* has failed to start" during the migration process, because although the source VM is powered on and running the P2V client, the guest operating system does not start.
- The time required for the migration depends on the size and number of volumes on the source system as well as the network bandwidth between the source and the target system. For example, transferring a source system with one 20 GB boot volume over a 1Gb network may take about 30 minutes.
- You can migrate multiple PMs or VMs at one time, but sharing network bandwidth may increase migration times.

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

Prerequisite: 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

1. Download the P2V client ISO file from the **Downloads** page at https://www.penguinsolutions.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.
2. If you want to verify the integrity of the ISO image, use the MD5 checksum hash function.

   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

1. Download the P2V client ISO file from the **Drivers and Tools** section of the **Downloads** page at https://www.penguinsolutions.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.
2. If you want to verify the integrity of the ISO image, use the MD5 checksum hash function.

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

   CertUtil -hashfile *path_to_file* MD5

   The **CertUtil** command displays a message indicating whether or not it completed successfully. If the command succeeds, continue with the next step. If the command fails, repeat the download.
3. Insert or connect the P2V client ISO file to the source VM and set the virtual CD drive as the boot device in the associated hypervisor.
4. Shut down the VM in preparation to boot the P2V client.

To migrate a PM or VM to the everRun system

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

5. 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 everRun system.)
6. The **# vCPUs** and **Memory(MB)** values are automatically detected and completed, but optionally modify them if you want the VM on the everRun system to have more CPUs or memory than the source PM or VM.
7. Specify the **Virt-v2v output options** for the target VM, as follows:
    a. Next to **Output to**, select **HA** (High Availability) or **FT** (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.)
8. If you want to save debugging messages from the migration process, optionally select the **Enable server-side debugging** check box. (The debugging messages are included if you generate a diagnostic file for your authorized Penguin Solutions 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.

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

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

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

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

To complete the migration on the everRun system

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

   Click **Finish** on the last wizard page to implement the changes.
3. Click **Start** to boot the new VM.
4. Click **Console** to open the console of the VM and log on to the guest operating system. (For information about using the console, see Opening a Virtual Machine Console Session.)

5. Disable any guest operating system services that are unnecessary for operation on the 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.

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 Penguin Solutions service representative, as described in Creating a Diagnostic File. The diagnostic file includes any server-side debugging messages from the migration process.

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

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

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

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

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

6. Open **Administrator Power Shell** and execute the following command:

   > **powercfg /h off**

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

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

A Linux-based VM may fail to exit the **booting** state in 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 (`/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, ). 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 (), the driver is working properly.

Related Topics

Adding, Migrating, and Importing Virtual Machines

Configuring Windows-based Virtual Machines

Configuring Linux-based Virtual Machines

Managing Virtual Machine Resources

Managing the Operation of a Virtual Machine

Importing Virtualization Files

To transfer a Virtual Machine (VM) from one system to another, or transfer an image that you created back to the same system to restore or duplicate the original VM, you can import an Open Virtualization Format (OVF) or an Open Virtual Appliance (or Application) (OVA) file as described in this section.

To migrate a physical machine (PM) or virtual machine (VM) to a system without using an OVF or OVA 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. You can import a VM from a remote source over the network.

  > **Note**: Instead of importing VM files from a remote source, you can pack your virtualization files into an ISO, upload the ISO to the system as a Virtual CD (VCD), and import the VM using the REST API. You can only import VMs from VCDs through the REST API — not through the everRun Availability Console. See Importing Virtual Machines from a VCD.

- 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 or OVA file from a local computer, a USB device, or a remote file system such as an NFS export or a Windows share (also known as a CIFS share, such as, for example, Samba). If you want to restore an existing VM on the same system to overwrite the VM and recover it from a previous backup copy, see Replacing/Restoring a Virtual Machine from an OVF File.

Notes:

- Import a VM if you are trying to create or clone a VM from a golden image, since the system will assign unique hardware ID and MAC addresses when importing a VM. (A golden image is typically a template VM created for the purpose of copying multiple times.) To prevent conflicts with the source VM, the **Import** wizard automatically assigns a new MAC address to each network interface in the new VM; however, you may need to manually update any IP addresses and host names as needed.
- You can import VMs only if they are running supported guest operating systems and boot interfaces, as described in Tested Guest Operating Systems.

  When you import a VM, the system imports the boot interface setting (BIOS or UEFI) from the OVF or OVA file; you cannot modify this setting.
- You can import a VM from a VMware source only if the source is running VMware Release 6.x.

  When importing a VMware VM, you must shut down the VM using operating system shutdown commands in addition to powering it off from the VMware console. If you shut down the VM using only the VMware console, the import will fail.
- If you import a VM from a VMware OVA file, ensure that your system has sufficient disk space for the operation. The system requires an amount of disk space equal to the size of the OVA file + the total size of the VM volume(s) to be created + 100 GB disk space that is temporarily reserved for expanding and processing the compressed OVA file. For example, if you need to import a 3 GB OVA file for a VM that requires a 32 GB volume, the minimum storage needed is 3 GB + 32 GB + 100GB = 135 GB.

  You can check the amount of **Free** disk space on your system on the **System** page of the everRun Availability Console under **Storage Allocation**. If your system lacks the amount of disk space needed to import a VMware OVA file, you can clear some disk space or instead migrate the VM directly over the network (with no OVF or OVA file) as described in Migrating a Physical Machine or Virtual Machine to a System.
- 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.

- If you begin to import an OVA file and then the node is placed into maintenance mode or loses power, the OVA import fails, and any future attempt to import an OVA file fails. For information on a work-around for this problem, see KB0014856.
- 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.
- After Importing a Linux VMware OVA file, you need to manually configure network information. See After Importing a Linux VMware OVA File, Manually Configure Network Information.

Prerequisite:

If you plan to import a VM image from an **OVF** file, you must first 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. You can also 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.

Prerequisite:
If you plan to import a VM image from an **OVA** file, you must first create the OVA file on a VMware system. The everRun system supports VMware OVA files that contain a metadata file and one or more disk image files.

To import an OVF or OVA file

1. Log on to the everRun Availability Console on the target everRun system.
2. 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 Virtual Machines Page), click **Import/Restore/Recover** to open the **Import/Restore/Recover Virtual Machine** wizard.
4. Select **Import/Restore**, and then click **Next**.
5. Select one of the following:
   - **Import from my PC**–Imports the VM from the PC running everRun Availability Console.

     Note: Browsing for VMware OVFs and OVAs is not supported when importing from a PC, but you can use any of the remaining methods to import VMware OVFs and OVAs.

     Click **Next** and then click **Browse** to locate the appropriate file on a local computer.
   - **Import from USB**–Imports the VM from a USB device mounted on the everRun system.

     Click **Next** and then select a partition from the pull-down menu. Click **List OVFs/OVAs** and select the appropriate file from the pull-down menu. You can optionally search for a file by entering the file name or partial file name in the *Search Files* box. The box lists OVA files that have names matching the name entered in the box, and that reside in various directories:
     - With the parent (root) directory as the search directory, the listed files reside in sub-directories in addition to the parent (root) directory.
     - With a sub-directory as the search directory, the listed files reside in the parent (root) directory in addition to the sub-directory.

- **Import from remote/network Windows Share (CIFS/SMB)**—Imports the VM from a Windows share on your local network. Note that the maximum length of the path to the VM, including the VM name, is 4096 characters.

  Click **Next** and 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/OVAs** and select the appropriate file from the list.

- **Import from remote/network NFS**—Imports the VM from an NFS share on your local network. Note that the maximum length of the path to the VM, including the VM name, is 4096 characters.

  Click **Next** and for **Repository**, enter the URL of the remote system in the format ***nnn.nnn.nnn.nnn/folder_name*** (do not include **http://** or **https://**).

  Click **List OVFs/OVAs** to display a list of all files in the remote folder. Select the appropriate file to import. You can optionally search for a file by entering the file name or partial file name in the *Search Files* box, or you can reorganize the list by clicking a column heading (*Name*, *Date Modified*, or *Size*). Click the file name to select the file, and then click **Next**.

If you have selected an OVA file, continue with the next step (import is the only option with an OVA file).

If you have selected an OVF file, click **Next**. Messages appear confirming whether or not it is a everRun-created file and whether or not you have the option to import or restore the VM. When selecting a everRun-created OVF file, you have the option of importing or restoring the file, and you can optionally display the following message:

> Restoring a VM attempts to preserve the hardware ID and MAC addresses of all network interfaces. Select **Restore** only if you are specifically trying to restore a particular instance of a VM and that it will be the only copy of this VM running across all systems on your network. Typically a **Restore** is used to recover a VM from a previous backup. Select **Import** if you are trying to create or clone a VM from a "golden" image, as this will assign a unique hardware ID and MAC addresses.

6. Select **Import** (scroll down the window, if necessary). (For a everRun-created OVF, you can also select **Restore**. See Replacing/Restoring a Virtual Machine from an OVF File for information.)
7. The wizard displays the **Prepare for Importing Virtual Machine** window, prompting you to upload additional files, if necessary. If prompted, select the appropriate file(s) to include for each volume

associated with the VM.

8. If you have selected an OVF file, you can review and, if necessary, edit the information (you may need to scroll down the window):

   - **Name, Boot Interface, CPU, and Memory**

     Displays the name of the VM, the boot interface, the number of vCPUs, and the total memory the VM can use. Edit the information, if necessary. (You cannot modify the **Boot Interface**; the system imports this setting from the OVF or OVA file.)

   - **Protection Level**

     Select the level of protection to use for the VM:

     - **Fault Tolerant (FT)**
     - **High Availability (HA)**
     - **Unprotected (NONE)**

     For information about these levels of protection, see Creating a New Virtual Machine and Modes of Operation.

   - **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.

The total number of networks cannot exceed the number of business networks on the everRun system. If you import the VM from an OVF file, you can select which networks to remove in the wizard. If you import the VM from an OVA file, the system automatically ignores the excess networks during the import process. In either case, you can connect more business networks to the everRun system before or after importing the VM to restore the network connections.

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

    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.

11. 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.
12. Click **Console** to open the console of the VM and log on to the guest operating system.
13. For Windows-based VMs only, download and update the VirtIO drivers to the latest supported versions, as described in Updating the VirtIO Drivers (Windows-based VMs). (The correct VirtIO drivers are already present in Linux-based VMs.)

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

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

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

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, ). 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 (), the driver is working properly.

After Importing a Linux VMware OVA File, Manually Configure Network Information

Importing a Linux VMware OVA file changes the network interface and `networks-scripts` file. After you import the file, you need to manually configure the network information using the following procedure:

1. On the **Virtual Machines** page, select the VM.
2. Click **Console** in the bottom panel to open the VM login page (for additional information, see Opening a Virtual Machine Console Session).
3. Log in to the VM.
4. Open a command prompt window.
5. Issue the `ifconfig` command. In the command output, check if `ip address` is assigned to the virtual network interface `eth0`.
6. If `ip address` is not assigned to `eth0`, list the contents of the `/etc/sysconfig/network-scripts` directory.
7. Note the value of `ifcfg-`*xxxx* (though not `ifcfg-lo`).

8. Rename `ifcfg-xxxx` to be `ifcfg-eth0`.
9. Edit the `ifcfg-eth0` file, changing the values of `DEVICE` and `ONBOOT`, as follows:

   ```
   DEVICE=eth0
   ONBOOT=yes
   ```

   Save the file.
10. Issue the following command to restart network services:

    ```
    systemctl restart network
    ```
11. Verify the IP assignment by issuing the command `ifconfig`. In the command output, confirm that `ip address` is assigned to `eth0`.

Related Topics

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

Adding, Migrating, and Importing Virtual Machines

Configuring Windows-based Virtual Machines

Configuring Linux-based Virtual Machines

Managing Virtual Machine Resources

Managing the Operation of a Virtual Machine

Importing Virtual Machines from a VCD

In cases where network bandwidth is insufficient for remote Virtual Machine (VM) import, you can pack virtualization files into an ISO, upload the ISO to the system as a Virtual CD (VCD), and import the VM using the Representational State Transfer (REST) application programming interface (API) as described in this section.



Note: You cannot import VMs from VCDs using the everRun Availability Console; you can import VMs from VCDs only through the REST API.

Prerequisites:

- Virtualization files – either an Open Virtualization Appliance (OVA) file, or an Open Virtualization Format (OVF) file with its associated disk files (VMDK or VHD)
- Ability to create an ISO
- Access to the everRun Availability Console in order to upload the VCD
- Access to the REST API in order to trigger the import

To import a VM from a VCD

1. Package your virtualization files into an ISO. For example, you can use the `genisoimage` CLI tool (part of the `cdrkit` suite) with the following command format:

   ```
   > $ genisoimage -o <output-file-name.iso> -allow-limited-size
   <space-separated virtualization and disk files>
   ```

   The following example creates an ISO image called `centos76-vhd.iso` from an OVF file and VHD disk file:

   ```
   > $ genisoimage -o centos76-vhd.iso -allow-limited-size
   CENTOS76.ovf disk0.vhd
   ```

   Note: Pack only one set of virtualization files per ISO (either one OVA file, or one OVF file with its disk files – not both).

   Note: Because the created ISO contains virtualization files (not a bootable image), you cannot use it for VM creation, only for VM import.

2. Upload the ISO as a Virtual CD:
   a. Log in to the everRun Availability Console.
   b. Open the **Virtual CDs** page and upload the ISO file (see Downloading a Virtual CD).
   c. To display a list of the files inside an uploaded VCD (optional):
      i. `ssh` into the system.
      ii. Run the command `blkid`. Look for the entry that matches the VCD name.

iii. List the ISO contents using the following command:

```
$ isoinfo -i /dev/mapper/<your-vcd-device-path> -l
```

d. Wait for the everRun system to synchronize the VCD across all nodes.

3. Import the VM using the REST API. You can submit the import request using any HTTP client, such as Swagger™, which is a suite of tools for the OpenAPI Specification. For more information, see REST API.

   For example, to execute the API call:

   a. Open Swagger at `https://<system IP or domain>:8888/restapi/swagger-ui/index.html#/`.

   b. Navigate to the section **VMImport Api**, then to `/v2/import/importvm`. This is the end-point to import a VM from an OVF or OVA file.

   c. Click **Try it out** and submit the request with the following required parameters:

      - `vm-name` – Name to assign to the imported VM.
      - `remote-type` – Type of remote connection or resource. Must be `vcd` for VCD-based import.
      - `vcd-name` – Name of the VCD you uploaded.
      - `network` – Network name.

      The following example includes all required parameters for a VCD-based OVF file import. Although certain parameters (such as `file-name` and `remote-dir`) are ignored for VCD imports, the request schema requires them, so placeholder values must be provided. See the Swagger documentation for details.

```
{
 "vm-name": "centos76-vhd",
 "file-name": "string",
 "remote-type": "vcd",
 "port": "string",
 "remote-dir": "string",
 "vcd-name": "centos76-vhd",
 "virtual-cpus": 0,
 "memory": 0,
```

```
 "destination": "string",
 "availability-level": "HA",
 "preferred-node": "string",
 "network": ["network0"]
}
```

After you submit the request, the system automatically:

- Resolves the VCD name to its internal block device path
- Mounts the VCD in read-only mode
- Locates and validates the virtualization files inside
- Parses the OVF or OVA file and imports the VM
- Cleans up (unmounts and removes temporary directories)

Related Topics

Adding, Migrating, and Importing Virtual Machines

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 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 Adding, Migrating, and Importing 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. However, the hardware ID 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.

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:

- 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 also cannot restore a VM from an OVA file.
- You typically restore a VM to replace a 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 everRun VM and that the restored VM will be the only copy of this VM running across all everRun 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.
2. 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 Virtual Machines Page), select the VM that you want to restore in the upper panel.
4. Do one of the following.
   - In the lower panel, click **Restore**.
   - Near the top pane, click **Import/Restore/Recover**, select **Import/Restore**, and click **Next**.
5. Select one of the following:
   - **Import from my PC**–Imports the VM from the PC running everRun Availability Console.
     a. Click **Next**.
     b. Click **Browse** to locate the appropriate folder on a local computer.
     c. Click the name of the desired file.
     d. Click **Open**.
   - **Import from USB**–Imports the VM from a USB device mounted on the everRun system.

Click **Next** and then select a partition from the pull-down menu. Click **List OVFs/OVAs** and select the appropriate file from the pull-down menu. You can optionally search for a file by entering the file name or partial file name in the *Search Files* box. The box lists OVA files that have names matching the name entered in the box, and that reside in various directories:

- With the parent (root) directory as the search directory, the listed files reside in sub-directories in addition to the parent (root) directory.
- With a sub-directory as the search directory, the listed files reside in the parent (root) directory in addition to the sub-directory.

- **Import from remote/network Windows Share(CIFS/SMB)**—Imports the VM from a Windows share on your local network.

  Click **Next** and 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/OVAs** and select the appropriate OVF file from the list.

- **Import from remote/network NFS**—Imports the VM from an NFS share on your local network.

  Click **Next** and for **Repository**, enter the URL of the remote system in the format ***nnn.nnn.nnn.nnn/folder_name*** (do not include **http://** or **https://**).

  Click **List OVFs/OVAs** to display a list of all files in the remote folder. Select the appropriate OVF file. You can optionally search for a file by entering the file name or partial file name in the *Search Files* box, or you can reorganize the list by clicking a column heading (*Name*, *Date Modified*, or *Size*). Click the file name to select the file, and then click **Next**.

6. Select **Restore**. (Scroll down the window, if necessary.) A warning message appears, stating that **Restore** will overwrite all existing data and configuration details and that you should proceed with caution.
7. Click **Continue** to proceed.
8. If prompted, add VHD files.
9. Review the information and make any desired edits, if necessary:

   - **Name, Boot Interface, CPU, and Memory**

     Displays the name of the VM, the boot interface, the number of vCPUs, and the total memory the VM can use. Edit the information, if necessary. (You cannot modify the **Boot Interface**; the system imports this setting from the OVF file.)

- **Protection Level**

  Select the level of protection to use for the VM:

  - **Fault Tolerant (FT)**
  - **High Availability (HA)**
  - **Unprotected (NONE)**

  For information about these levels of protection, see Creating a New Virtual Machine and Modes of Operation.

- **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.

- **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.

  The total number of networks cannot exceed the number of business networks on the everRun system. You can select which networks to remove in the wizard, or connect more business networks to the everRun system before or after restoring the VM to restore the network connections.

10. 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

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

Adding, Migrating, and Importing 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. Alternatively, you can back up a running VM to a set of XML and raw hard disk data files, as described in Backing Up a Running Virtual Machine.

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:

- Because the source VM must be shut down to export it, 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 VM volume data that you can export to an external file share or USB device formatted as a FAT or VFAT file system is 4 GB, even if the total size of the target device is much larger. If you attempt to export data more than 4 GB to a FAT or VFAT file system, the export will fail.
- For Linux-based VMs, when exporting a VM to another system, you do not need to modify the `/etc/fstab` file.
- For Ubuntu-based VMs running some older Ubuntu releases, you may need to edit the `/boot/grub/grub.cfg` file and change the `gfxmode` parameter to `text` (for example, `set gfxmode=text`) before exporting a VM; otherwise, the new VM's console may hang on another system. You can restore the original setting in the source VM after the migration.

Prerequisites:

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

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

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

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

1. Log on to the everRun system with the everRun Availability Console.
2. On the **Virtual Machines** page, select the VM to export.
3. Click **Console** to open the console of the VM and log on to the Windows guest operating system.
4. Ensure that all volumes are labeled accurately, as summarized in Managing Windows Drive Labels.
5. If you plan to create a golden image of a VM that you can export to other systems for redeployment, run the Windows System Preparation Tool (`Sysprep`) to prepare the guest operating system for redeployment.

> **Caution**: Do not use `Sysprep` on a VM that is needed for production use, because this utility resets the configuration of the guest operating system to its default settings for redeployment. If needed, you can make a copy of the VM, run `Sysprep` on the copy, and then export the copy.

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 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.nnn.nnn.nnn*** (do not include **http://** or **https://**).

   - **Mount USB**

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

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

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

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

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

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

The 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 Virtualization Files.

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

Adding, Migrating, and Importing 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.

You can also mount a USB-attached SCSI (UAS) compliant device, and then use it to import to as well as restore or export from, in the same way that you can use other USB devices.

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

Notes:

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

Prerequisite: Prepare the mount location:

- If you are using a USB device to export or import a VM, attach the device to the current primary node for the system (displayed as **node*n* (primary)** on the **Physical Machines** page). Confirm that the system displays the USB device: navigate to the **Physical Machines** page, click the node to which you attached the device, and in the lower pane, select the **USB Device** tab. The USB device you attached should appear in the tab's display.
- If you are using a network-mounted folder for a Windows/CIFS share or an NFS export, create the folder in your environment where you can store the exported VM or 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 mounting an NFS export of CIFS share.

To mount a USB device or network-mounted folder

1. On the **Virtual Machines** page, select a VM, 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.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: If you plan to run `Sysprep` to create a golden image of a VM that you can export to other systems for redeployment, ensure that each volume has a unique identifiable label before running `Sysprep`. 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

Adding, Migrating, and Importing Virtual Machines

Configuring Windows-based Virtual Machines

Understanding Windows 11 Virtual Machine Differences from Other VMs

Windows 11 imposes stricter requirements for security than other operating systems. Supporting a Windows 11 virtual machine (VM) requires several emulated components and configurations to enable these security features. everRun automatically makes these changes in any Windows 11 VM, including:

- Enabling Secure boot

  This also requires the system to change the chipset emulation from i440fx to q35, and requires the user to select the **UEFI boot interface** when creating the VM.
- Configuring virtual Trusted Module Platform (vTPM) security
- Setting the vCPU model for the VM to be newer than Sandy Bridge

Notes:

- These settings cannot be modified, and they are not available to any other guest operating system.
- Because secure boot requires the UEFI boot interface, attempting to create a Windows 11 VM with the BIOS boot interface will cause an error. To correct the issue, you need to delete the VM and create a new VM with the UEFI boot interface.

To verify that your PMs can support running Windows 11 VMs

Open the **Virtual Machines** page of the everRun Availability Console. In the upper-right corner, next to the **Create** button, check the status of the Windows 11 compatibility icon:

-  Windows 11: The PMs support running Windows 11 VMs.
-  Windows 11: The PMs do not support running Windows 11 VMs.

To configure the UEFI Boot interface for the VM

When running the **VM Creation** wizard, under **Boot Interface**, select **UEFI**.

To see which vCPU processor type is being presented to the VM

In the Windows 11 guest operating system, open the **Control Panel**, click **System**, and click **About**.

To verify that secure boot and TPM are enabled on the VM

In the Windows 11 guest operating system, right-click the **Start** button and click **Run** to open the **Run** window, type **tpm.msc**, and click **OK**. The **Trusted Platform Module (TPM) Management** snap-in appears. The **Status** should be similar to **The TPM is on and ownership has been taken**.

For more information about TPM, see the Microsoft Windows documentation.

Related Topics

Adding, Migrating, and Importing Virtual Machines

Managing Virtual Machine Resources

Managing the Operation of a Virtual Machine

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).

A VCD with the ISO file of the VirtIO drivers is installed on the system during the installation of system software as well as during an upgrade of the system software. To confirm that the VCD exists, check **The Virtual CDs Page** (see the Virtual CDs Page) for a VCD with **virtio** in the name. If the VCD exists, update the VirtIO drivers (see To update the VirtIO drivers in a Windows-based virtual machine). If the VCD does not exist, create it (see To download the VirtIO drivers and create a VCD) and then update the drivers.

Notes:

- 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.penguinsolutions.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 VirtIO Driver Update**.
    c. Click the link to the appropriate file.

    Ensure that you download the version of the VirtIO ISO file that matches the version of your everRun system.

2. If you want to verify the integrity of the ISO image, use the MD5 checksum hash function.

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

   CertUtil -hashfile ***path_to_file*** **MD5**

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

3. Open the 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.

5. 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

Adding, Migrating, and Importing 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

Opening a Virtual Machine Console Session

Configuring Windows-based Virtual Machines

Adding, Migrating, and Importing 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

Adding, Migrating, and Importing 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:

- 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.penguinsolutions.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

Adding, Migrating, and Importing 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 everRun Availability Console, edit the `/boot/grub/grub.cfg` file and change the `gfxmode` parameter to `text` (for example, `set gfxmode=text`). If the VM console hangs before you can set the parameter, see the troubleshooting information in Opening a Virtual Machine Console Session 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

Adding, Migrating, and Importing 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

Adding, Migrating, and Importing 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 transaction-intensive applications. In this case, taking a snapshot without preparing your applications results in a *crash-consistent* snapshot, as if you took the snapshot after a power outage.

If your applications support QEMU signaling, the everRun software can signal your applications through the QEMU guest agent to ensure that your applications *quiesce* or freeze during a snapshot on your everRun system, thus ensuring an application-consistent snapshot.

Most Linux distributions already include a QEMU guest agent (usually in the `qemu-guest-agent` package). For information about installing and configuring the QEMU guest agent, see the documentation for your particular Linux distribution.

Caution: Before installing a QEMU guest agent, contact your application vendor(s) to determine if they support QEMU signaling and if any additional configuration steps are necessary to quiesce your applications. You can create application-consistent snapshots only if your applications support QEMU signaling and the QEMU guest agent is properly installed and running.

Notes:

- By default, all snapshots are considered crash-consistent snapshots unless you explicitly install the QEMU guest agent and configure your applications to quiesce when signaled by the everRun software.
- When you install the QEMU guest agent, you may need to restart your VMs. If your VMs are use, schedule a maintenance period for the installation.

Related Topics

Configuring Linux-based Virtual Machines

Adding, Migrating, and Importing 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

To back up a running virtual machine, see Backing Up a Running 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

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 the 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** () in the bottom panel.

Note:
After you click **Console**, the console session that opens may be blank if the browser has an HTTPS connection to the system, but does not have a security exception for it. In this situation, click the IP address in the upper-right corner of the session window. This IP address, which is in the format https://*system_IP_address*:8000, adds the system IP address as a security exception site in the browser. A security exception allows the browser to open the site.

Depending on the browser, additional security windows or messages may appear. With some browsers, one or more security messages appear, and you need to click through those messages. With other browsers, the address bar turns red with no message, and you need to click the address to proceed. Some specific examples are:

- If **Certificate error** appears in the address bar, you may need to (1) click the address; (2) on a page displaying **The website cannot display the page**, click **More information**; and then (3) on a page displaying **This site is not secure**, click **Go on to the webpage (not recommended)**.
- If the page **Warning: Potential Security Risk Ahead** appears, click **Advanced** and in the next window, click **Accept Risk and Continue**.
- If **Error response** with **Error code 405** appears, close the window or tab.

This security exception will then apply to all VMs. You need to perform these actions only once for each browser. When you click **Console** in the future, the console session to the VM opens successfully.

After you have opened the VM console session, you can resize the browser window and the VM console session. You can also use keyboard shortcuts.

To resize the browser window and the VM session

1. Open the VM console session (see procedure above).

   Icons appear at the left edge of the window. To display the icons, you may need to click the arrow in the tab at the left edge of the window.

2. To resize the browser window to full screen, click the full-screen icon ().

   When in full screen, click the full-screen icon () again to resize the browser to a smaller window.

3. To resize the VM session inside the browser, click the Settings icon ( ) and select a **Scaling Mode** (click the current mode to view a pull-down menu with other settings):
   - **Remote Resizing** (the default)—The size of the VM session changes when you change the resolution of the guest OS.
   - **Local Scaling**—The size of the VM session changes automatically to fill the full screen with the original width and height ratio.

To use keyboard shortcuts

1. Open the VM console session (see procedure above).

   Icons appear at the left edge of the window. To display the icons, you may need to click the arrow in the tab at the left edge of the window.
2. Click the **A** icon ( ) at the left edge of the window to display the keyboard shortcut-selection icons.
3. The following icons appear:
   - Ctrl—Click for the **Ctrl**-key function.
   - Alt—Click for the **Alt**-key function.
   - —Click for the **Tab**-key function.
   - Esc—Click for the **Esc**-key function.
   - —Click for the **Ctrl**+**Alt**+**Delete**-keys function.

Troubleshooting

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

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

To resolve an issue where the VM console window is blank

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

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

Close all console windows and open only one console window.

To resolve an issue where the VM console window hangs on the everRun system

For Ubuntu-based VMs, the VM console hangs in the everRun Availability Console if you do not properly set the `gfxmode` parameter. In the guest operating system, edit the `/boot/grub/grub.cfg` file and change the `gfxmode` parameter to `text` (for example, `set gfxmode=text`).

If the console hangs before you can set the parameter, do the following:

1. Restart the VM in the everRun Availability Console.
2. At the GRUB menu, press `e` to edit the grub command.
3. On the next screen, on the `gfxmode` line, change `$linux_gfx_mode` to `text` so the line reads:

   ```
   gfxmode text
   ```

4. Press **Ctrl-x** or **F10** to boot the guest operating system.
5. To update the setting so it persists for each boot cycle, edit the `/boot/grub/grub.cfg` file and change the `gfxmode` parameter to `text` so the line reads:

   ```
   set gfxmode=text
   ```

6. Save the `/boot/grub/grub.cfg` file.

To change the terminal type in a Linux-based VM if the console screen is unreadable

By default, the Linux operating system sets the `TERM` variable to `vt100-nav`, which is not properly supported by the `vncterm` program, the basis for the VM console in everRun Availability Console. If you use anything other than the command line, the screen becomes unreadable. To resolve this issue, change the terminal type in the Linux guest operating system:

1. Open the `inittab` file in the guest operating system.
2. In the following line, replace `vt100-nav` with `vt100` by deleting `-nav` at the end of the line. The updated line appears as follows:

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

3. Save the **inittab** file.

Related Topics

Starting a Virtual Machine

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

Adding, Migrating, and Importing 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:

- 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

Adding, Migrating, and Importing 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 or change its vCPU emulation, 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 number of virtual CPUs (vCPUs), vCPU emulation setting, memory allocation, storage resources, 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.

Note: When reprovisioning a VM, you cannot change the protection level or preferred node of an unprotected VM. For more information, see Unprotected Operation.

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:
    a. Type the **Name** and an optional **Description** for the VM as you want them to appear in the everRun Availability Console. For VM name requirements, see Creating a New Virtual Machine.
    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:
    a. In the **vCPUs** field, specify the number of vCPUs.
    b. Click **Advanced options** to expand the vCPU emulation section. Select the **vCPU Emulation** setting for this VM from the pulldown menu. Available options are Skylake, Haswell, and Sandybridge. The new setting persists across VM shutdown, restart, and upgrade, unless you change it again.
    c. In the **Memory** section, set the amount of memory to assign to the VM. For more information, see Planning Virtual Machine vCPUs and Planning Virtual Machine Memory.
    d. Click **Next**.

6. On the **Volumes** page, you can:

> **Notes**:
>
> - 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 **Volume 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**.

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

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

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

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

8. On the **Configuration Summary** page:

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

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

   b. To accept the VM as provisioned, click **Finish**.

9. Click **Start** to restart the VM.

10. For Windows-based VMs, if you changed the number of assigned virtual CPUs in a Windows-based VM from 1 to *n* or *n* to 1, after restarting the VM at the end of the re-provisioning process, you must shut down and restart the VM a second time. This allows the VM to correctly reconfigure itself for Symmetric Multiprocessing (SMP). The VM displays odd behavior and is not usable until it is restarted.

Related Topics

Managing Virtual Machine Resources

Planning Virtual Machine Resources

Managing Virtual Machines

Creating a Volume in a Virtual Machine

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

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

To create a new volume in a VM

1. Open the **Virtual Machines** page (see the Virtual Machines Page).
2. Select a VM and click **Shutdown**.
3. When the VM has stopped, click **Config** to display the **Reprovision Virtual Machine** wizard.
4. Click **Next** on each wizard page until the **Volumes** page is displayed. (If applicable, see Reprovisioning Virtual Machine Resources to configure additional VM resources.)
5. On the **Volumes** page, click **Add a new volume**. (If the button is not visible, scroll down to the bottom of the wizard page.)
6. Under **To Be Created**, do the following:
    a. Type the **Name** of the volume as it will appear in the everRun Availability Console.

    The volume name must meet the following requirements:

    - Valid characters are a-z, A-Z, 0-9 , +, _, and . (dot).
    - A volume name cannot be only . or .. (dots).
    - A volume name cannot be snapshot or pvmove.
    - A volume name cannot contain any of the following reserved strings: _cdata, _cmeta, _corig, _mlog, _mimage, _pmspare, _rimage, _rmeta, _tdata, _tmeta, or _vorigin.

    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

Attach a volume to connect a currently unused volume to a virtual machine.

Note: If you attach a boot volume to a VM that already has a boot volume, the newly added volume is attached as a data volume. You might want to attach a volume in this manner to diagnose a boot problem or data corruption in another VM's boot volume. After using guest operating system tools to resolve the issue, detach the volume and reattach it to its original VM.

Prerequisite: Before attaching a volume to a virtual machine, you must shut down the virtual machine.

To attach a volume to a virtual machine

1. Ensure that the volume you want to attach is not in use by another virtual machine; otherwise, you cannot attach it. Open the **Volumes** page, locate the volume, and ensure that the value in the **Used**

By column is **None**.

2. Open the **Virtual Machines** page (see the Virtual Machines Page).
3. Select a VM and click **Shutdown**.
4. When the VM has stopped, click **Config** to display the **Reprovision Virtual Machine** wizard.
5. Click **Next** on each wizard page until the **Volumes** page is displayed. (If applicable, see Reprovisioning Virtual Machine Resources to configure additional VM resources.)
6. On the **Volumes** page, locate the pulldown menu next to the **Add a new volume** button. Select an unattached volume from the pulldown menu and click **Attach**.

   (If the pulldown menu is not visible, scroll down to the bottom of the wizard page. The pulldown menu is displayed only if there are unattached volumes on the 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:

- 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.

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

7. Click **Next** on each wizard page until the **Configuration Summary** page is displayed. Verify the

configuration changes.

8. Click **Finish** to detach the selected volume.

Related Topics

Attaching a Volume to a Virtual Machine

Removing a Volume from a Virtual Machine

Managing Virtual Machine Resources

Planning Virtual Machine Resources

Managing Virtual Machines

Removing a Volume from a Virtual Machine

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

- 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.

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

To remove an unattached volume

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

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

Related Topics

Detaching a Volume from a Virtual Machine

Attaching a Volume to a Virtual Machine

Managing Virtual Machine Resources

Planning Virtual Machine Resources

Managing Virtual Machines

Renaming a Volume on the 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**. For volume naming restrictions, see Creating a Volume in a Virtual Machine.

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

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), select the volume that you want to expand.

3. In the bottom pane, click the **Container** tab and then click **Expand Container**.
4. 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

1. 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:

1. You cannot enable (or disable) the same component on both nodes at the same time.
2. You cannot disable both instances of a VM.
3. 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

A virtual CD (VCD) is a read-only ISO image file that resides on a storage device of the everRun system. You can use VCDs to make software installation media available to the VMs on your everRun system in ISO format.

You can also use VCDs to import a virtual machine (VM). See Importing Virtual Machines from a VCD.

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 download a VCD to your local computer. 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
- Inserting a Virtual CD
- Ejecting a Virtual CD
- Booting from a Virtual CD
- Renaming a Virtual CD
- Downloading 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.
2. If you create a bootable VCD for installation, it must be a single CD or DVD. Multiple CDs or DVDs are not supported.

To create a VCD

1. If necessary, create ISO files of any physical media for which you will create VCDs.
2. Open the **Virtual CDs** page in the everRun Availability Console.
3. Click **Create VCD** to open the **Virtual CD Creation** wizard.
4. In the wizard, select a storage group with sufficient free space for the VCD.
5. Type a name for the VCD.
6. Select a source for the VCD:
   - **Upload ISO file** uploads a file from your system running the everRun Availability Console. Click **Browse**, select the ISO file on your system, and click **Open**.
   - **Copy CD ISO from network source** copies the file from a Web URL. Specify the URL of the ISO file.
7. Click **Finish** to upload or copy the ISO file from the specified source.

   The **Virtual CD Creation** wizard displays progress of the upload.

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

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

Related Topics

Inserting a Virtual CD

Ejecting a Virtual CD

Managing Virtual CDs

Adding, Migrating, and Importing Virtual Machines

Inserting a Virtual CD

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

Caution: When you insert a VCD into a running, 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

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

Adding, Migrating, and Importing 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

Downloading a Virtual CD

Download a virtual CD (VCD) to make the software on the VCD available for uploading at a future time.

Prerequisite: You must first create a VCD, if you have not yet done so. See Creating a Virtual CD.

To download a VCD

1. Open the **Virtual CDs** page in the everRun Availability Console.
2. Click the name of the VCD you want to download.
3. Click **Download**. A window opens, displaying a folder on your local computer.
4. Select a destination for the file and click **Save**.

Depending on the size of the file, the download may require several minutes to complete.

Related Topics

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:

- 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

Adding, Migrating, and Importing 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:

- 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.
- You cannot create a snapshot of an unprotected VM. For more information, see Unprotected Operation.

Prerequisite: Both PMs of the everRun system must be online for the snapshot process to function properly. If only one PM is online, the snapshot is written only to the online PM, and that same PM must be primary if you export the snapshot later.

To prepare for creating a snapshot (Windows-based VMs only)

1. If you want to create an application-consistent snapshot, ensure that the QEMU guest agent is installed and running, as described in Installing the QEMU Guest Agent for Application-Consistent Snapshots (Windows-based VMs).
2. Ensure that all volumes are labeled accurately, as summarized in Managing Windows Drive Labels.

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 the VM for which you want to create a snapshot and ensure that it is in the **running** state and healthy. See the Virtual Machines Page.
4. With the VM selected, click the **Snapshot** button in the bottom pane.

   The **Snapshot Virtual Machine** dialog box appears, displaying the rows **Data Volumes to Capture** and **Boot Volume to Capture**. Each row includes the columns **Name**, **Type**, **Space Required**, **Space Available**, **Node0**, and **Node1**.
5. In the **Snapshot Virtual Machine** dialog box, all volumes 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.

   > **Note**: A broken icon (✖) in a node column indicates that the node has insufficient container space for the snapshot.

   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

Adding, Migrating, and Importing 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 Adding, Migrating, and Importing 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:

- 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 (✔) the **State** column. If a snapshot is broken (✖), it may indicate that one or more volumes in the snapshot are unavailable for the VM creation.
5. 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.
9. 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

Adding, Migrating, and Importing 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 Adding, Migrating, and Importing 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:

- 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 everRun 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.
2. 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 (✔) the **State** column. If a snapshot is broken (✖), it may indicate that one or more volumes in the snapshot are unavailable for export. You can check volume availability in step 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.nnn.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 Virtualization Files.

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

Adding, Migrating, and Importing 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 storage group.
- When you remove a snapshot, 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. So, if you need to delete a snapshot, do so while its associated VM is running, or allow the associated VM (if stopped) to remain stopped.
  - A user must **not** shutdown a VM associated with the coalescing snapshot(s). You must **not** shut down the associated VM from the guest operating system because doing so causes data corruption. You cannot shut down the VM using the everRun Availability Console because the console prevents you from doing so.
  - 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 KB0013465.

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**.
5. A confirmation window appears. Click **Yes** to continue deleting the snapshot or click **No** to stop the deletion.

Related Topics

Managing Snapshots

Adding, Migrating, and Importing 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:

1. By default, the everRun software automatically assigns MAC addresses to the VMs. Do not override the default settings unless you have specific requirements (for example, to support software applications that are licensed on a MAC-address basis).
2. If you change the **Static System IP** address, any MAC addresses automatically assigned to the VMs will change when the VMs reboot, because the everRun software generates MAC addresses for the VMs based on the system IP address. To prevent changes to the MAC address for a VM, set a persistent MAC address as described in the following procedure. Contact your network administrator to generate a valid MAC address for your environment, and remember to update any firewall rules based on the new MAC address.

Prerequisite: Before modifying the MAC address for a virtual machine, you must shut down the VM.

To assign a specific MAC address to a VM

1. Open the **Virtual Machines** page (see the Virtual Machines Page).
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

Advanced Topics (Virtual Machines)

Managing the Operation of a Virtual Machine

Forcing a VM to Boot

You can force a VM to boot using the **Force Boot** button on the VIRTUAL MACHINES page. However, the **Force Boot** button is active only when the 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 Penguin Solutions 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 green-checked copy of a different volume, and the volume's mirror copy on the opposite node is not green-checked.
- Both nodes are required because multiple VMs are degraded, yet are green-checked on opposite nodes (for example, Node0 has a good boot volume but a bad data volume, while Node1 has a bad boot volume but a good data volume).

If you perform a **Force Boot** on a system with outdated volumes, contact your authorized Penguin Solutions 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 Penguin Solutions 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.
- Contact your authorized Penguin Solutions service representative.

To force a VM to boot

After you have consulted with your authorized Penguin Solutions 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, click **Virtual Machines** in the left panel.
2. Navigate to the **Virtual Machines** page.
3. On the **Virtual Machines** page, select the VM that is stopped and that you want to force to boot (for example, VM-1).
4. In the bottom panel, click the **Start** button.

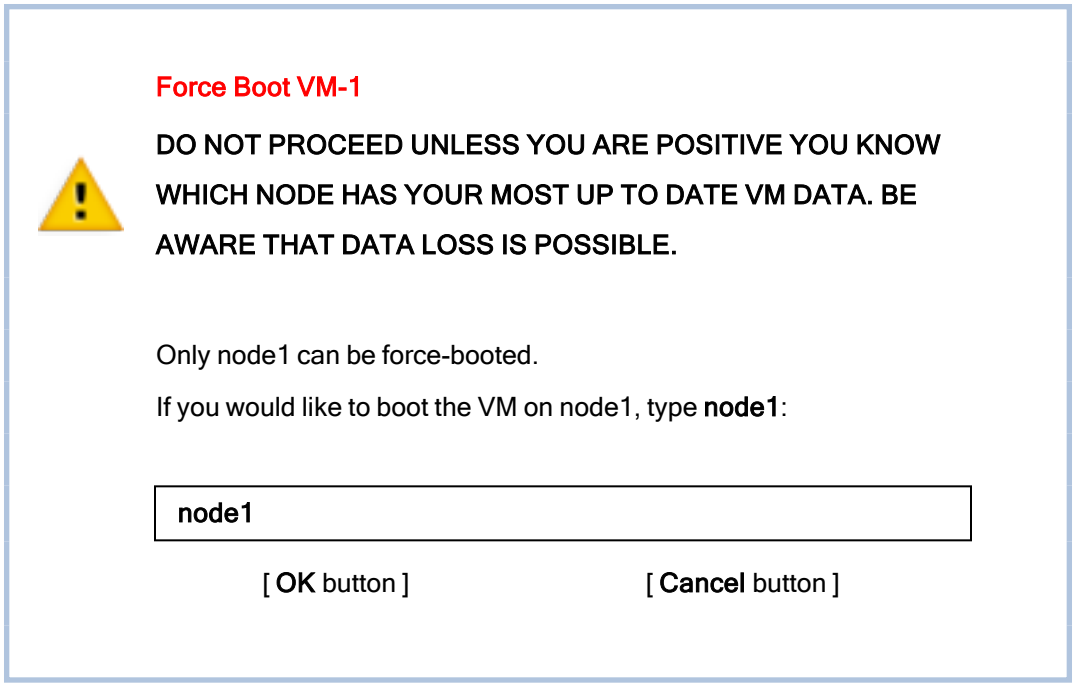
   The VM begins to boot. It continues booting until the time-out limit is reached, possibly as long as 5 minutes. When the time-out limit is reached, the **Force Boot** button becomes active.
5. To force the VM to boot, click **Force Boot**.

   A warning appears, asking you if you are positive that you know which node has the most up-to-date VM data. The warning also tells you to be aware that data loss is possible. In addition, a message tells you the node on which you can force the VM to boot.

   Caution: If you select the wrong node during **Force Boot**, data is damaged.

   You must type the node (node0 or node1) as indicated in the message. The following message is an example:



6. Click **OK** to force the node (for example, node1) to boot. (Click **Cancel** to cancel the procedure.) As the force-boot process begins and continues, additional confirmation messages appear before the VM starts and the data is marked as valid to the system.

   The VM begins to run. On the **Virtual Machines** page, the VM is listed with a warning because the node (for example, node0) is still offline.

   Once the secondary node is brought back in to the system, all data synchronizes from the node running the VM. In this example, all data synchronizes from node1 to node0.

Related Topics

Advanced Topics (Virtual Machines)

Managing the Operation of a Virtual Machine

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.

Note: You cannot change the protection level or preferred node of an unprotected VM. For more information, see Unprotected Operation.

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
App	2	30 secs
DB	2	10 mins
Web	3	0

1. everRun boots the DNS VM.
2. 2 minutes after the DNS VM is started, everRun starts the App and DB servers in group 2.
3. 10 minutes after the DB VM is started, everRun starts the Web VM in group 3.

- **OS and Application Start Time** should be set to the time it takes from starting the virtual machine until the guest operating system and applications are fully functional.

Related Topics

Advanced Topics (Virtual Machines)

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 Penguin Solutions service representative, as doing so may affect the continuous uptime of your system.

Notes:

1. The **Reset Device** button is displayed only if the virtual machine falls below its MBTF threshold.
2. The **Clear MTBF** button is displayed only if the system software supporting a VM on one physical machine falls below its MBTF threshold.

To reset the MTBF counter for a virtual machine

1. On the **Virtual Machines** page, select a virtual machine.
2. Click **Reset Device**.

If the system software supporting a VM on one physical machine fails too often, perform the steps below to reset its MTBF counter.

To reset the MTBF counter for a VM on one physical machine

1. On the **Virtual Machines** page, select a virtual machine.
2. Click **Clear MTBF**.

Related Topics

Advanced Topics (Virtual Machines)

Managing the Operation of a Virtual Machine

Creating a Diagnostic File

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.)

You can also attach a USB-attached SCSI (UAS) compliant device, and then use it to import to as well as restore or export from, in the same way that you can use other USB devices. In addition, you can access the UAS device from a Windows or Linux guest operating system. For details about accessing UAS devices from guests, see KB0015134.



Caution:
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.

Notes:

1. You can attach only supported USB devices to a guest operating system. everRun systems support up to and including USB 3.2 Gen 1 (5 Gbps) devices. everRun systems do not support USB 3.2 Gen 2 (10 Gbps), or higher, devices in the guest operating system. However, you can insert a Gen 2 or higher device into a Gen 1 host port, which forces the device to operate at the Gen 1 (5 Gbps) speed; in this case, you can attach the device to a guest operating system. (USB 3.2 Gen 1 (5 Gbps) devices are formerly referred to as USB 3.1 Gen 1 devices, and USB 3.2 Gen 2 (10 Gbps) devices are formerly referred to as USB 3.1 Gen 2 devices.)
2. Some USB devices claim to be 3.0 compatible, but are not. If you are using such a device, switch to another USB device that is 3.0 compatible.
3. The VM must be running in order for you to attach a USB device to it.
4. By default, USB devices are enabled for attachment to VMs. To change this configuration, see Configuring VM Devices.
5. Use either of the following methods to detach (that is, eject) a supported USB device from a Windows-based VM:
   - Clicking Eject in File Explorer—If you eject the device from File Explorer, you must detach it in the everRun Availability Console using the procedure below. Then, physically remove it from the everRun system and reinsert it before reattaching to the same or another VM.
   - Clicking Safely Remove Hardware and Eject Media in the taskbar—If you eject the device from the taskbar, you must detach it in the everRun Availability Console using the procedure below. You do not need to physically remove it from the everRun system before reattaching it to the same or another VM.

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

8

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

everRun Availability Console

Physical Machines and Virtual Machines

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.
3. Select the appropriate PM (node0 or node1) and then click **Work On**, which changes the PM's **Overall State** to **Maintenance Mode** and the **Activity** state to **running (in Maintenance)**.
4. Follow vendor instructions for adding or replacing a hot-swappable component in the PM.
5. Select the repaired PM on the **Physical Machines** page. Click **Finalize** and then click **OK**.

If you add a cable to both PMs and if they are on the same subnet, everRun detects the connectivity and pairs the NICs in a newly created shared network. You have the option of renaming the newly created shared network on the **network** page.

Related Topics

Maintenance Mode

Maintaining Physical Machines

everRun Availability Console

Physical Machines and Virtual Machines

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.
3. Select the appropriate PM (node0 or node1) and then click **Work On**, which changes the PM's **Overall State** to **Maintenance Mode** and the **Activity** state to **running (in Maintenance)**.
4. After the PM displays **running (in Maintenance)**, click **Shutdown** and then **OK**.
5. Add or replace the component.
6. If you disconnected any network cables, reconnect them. Do not add cables to any new network ports in this step.
7. On the PM that is shutdown, press the power button. As the PM powers on, everRun also powers on and begins synchronizing the PM's storage (appears).
8. On the **Networks** page, click the **Fix** button, if it is highlighted, which may occur when network cables have been moved on the upgraded PM.
9. Select the repaired PM on the **Physical Machines** page. Click **Finalize** and then click **OK**.
10. When synchronization is complete (disappears), perform steps 3 through 9 for the other PM, if necessary.

Note: To avoid data loss, do not power down the primary PM while the disks are synchronizing.

Related Topics

Maintenance Mode

Maintaining Physical Machines

everRun Availability Console

Physical Machines and Virtual Machines

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

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

Physical Machines Page

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 Penguin Solutions 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 need to replace the preferred node that contains an unprotected VM, you must shut down the unprotected VM, export the VM, and delete the VM before replacing the node; otherwise, you will be unable to boot the VM on the replacement node. After replacing the node, you can import the VM and configure it as an unprotected VM again. For more information about unprotected VMs, see Unprotected Operation.

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 Penguin Solutions 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.
4. Configure the firmware (BIOS or UEFI). See Configuring Settings in the Firmware Setup Utility.

> **Note**:
> You must re-activate 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.
2. Select the appropriate PM (node0 or node1) and then click **Work On**, which changes the PM's **Overall State** to **Maintenance Mode** and the **Activity** state to **running (in Maintenance)**.
3. After the PM displays **running (in Maintenance)**, click **Recover**.
4. When prompted to select the type of repair, click **PXE PM Replace - Initialize All Disks**.

   > **Caution**: Selecting **PXE PM Replace - Initialize All Disks** deletes all data on the replacement PM.

5. Select one of the following PXE Settings:

   - **Only respond to PXE requests from the current partner node.**

     Waits for a PXE boot request from the MAC address of the current partner node. Select this option if you are recovering the existing PM by completely wiping and reinstalling it (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 everRun 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.
7. 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:

    1. If you are not sure of which port to use, use the arrow keys to select one of the ports, and click the **Identify** button. The LED on the selected port will then flash for 30 seconds, allowing you to identify it. Since the LED may also flash due to activity on that network, you can 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.

    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.

14. 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.

    > **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.

19. To activate the replacement PM, re-activate the product license for the everRun system. On the **Preferences** page, click **Product License** and click **Check License Now** to automatically activate the license. If the system has no Internet access, 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.

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, open the **Physical Machines** page and 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.

2. Select the appropriate PM (node0 or node1) and then click **Work On**, which changes the PM's **Overall State** to **Maintenance Mode** and the **Activity** state to **running (in Maintenance)**.
3. After the PM displays **running (in Maintenance)**, click **Recover**.
4. When prompted to select the type of repair, click **DVD/USB PM Replace - Initialize All Disks**.

   Caution: Selecting **DVD/USB PM Replace - Initialize All Disks** deletes all data on the replacement PM.

5. Click **Continue** to begin the replacement process. The system shuts down the PM in preparation for the system software reinstallation.
6. After the PM is powered off, install the replacement PM 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.
8. 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**:
>
> 1. If you are not sure of which port to use, use the arrow keys to select one of the ports, and click the **Identify** button. The LED on the selected port will then flash for 30 seconds, allowing you to identify it. Since the LED may also flash due to activity on that network, you can 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 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.

15. 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.

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.

18. 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:

- 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.

20. To activate the replacement PM, re-activate the product license for the everRun system. On the **Preferences** page, click **Product License** and click **Check License Now** to automatically activate the license. If the system has no Internet access, 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, open the **Physical Machines** page and 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

everRun Availability Console

Physical Machines and Virtual Machines

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

1. 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

everRun Availability Console

Physical Machines and Virtual Machines

Physical Machines Page

Part 2: Supporting Documents

See the following support documents for release information, and reference and troubleshooting information.

- everRun Release 8.1.1.0 Release Notes
- everRun Command Line Interface Reference
- System Reference Information
- Security
- SNMP

9

Chapter 9: everRun Release 8.1.1.0 Release Notes

These Release Notes (updated at 1:20 AM on 10/29/2026) are for everRun Release 8.1.1.0.

See the following sections:

- New Features and Enhancements
- Features No Longer Supported
- Bug Fixes
- CVE Fixes
- Important Considerations
- Known Issues
- Accessing Stratus Knowledge Base Articles
- Getting Help

New Features and Enhancements

everRun Release 8.1.1.0 includes the following new features:

- **vCPU Emulation Enhancements** – Enhancements to Virtual Machine (VM)-related controls, UI messages, and alerts, which improve VM compatibility and predictability during software upgrades and processor changes. For more information, see Managing Virtual Machines.
- **Backup of Shut-Down VMs** – You can now back up VMs that are in a shut-down state, enabling more flexible disaster recovery and backup management without requiring VMs to be running.

- **Windows 11 and Virtual Trusted Platform Module (vTPM) Improvements** – Significant improvements to Windows 11 VM support, including reliable vTPM handling after power outages, correct vCPU emulation enforcement, and the ability to upgrade Windows 10 guests to Windows 11 (both BIOS and UEFI).
- **New Guest OS Support** – New guest operating system support for Red Hat Enterprise Linux 9.8, 10.0, 10.1, and 10.2 (see Tested Guest Operating Systems for full list).
- **Upgrade Process Improvements** – Support for Advanced Package Tool (apt) upgrade, the standard, Ubuntu-based upgrade process. Starting with Release 8.1.1.0, upgrade now uses apt-upgrade only. In the everRun Availability Console, add the kit and run the upgrade as described in Upgrading everRun Software Using an Upgrade Kit. On dual-node systems, one node is upgraded at a time so that VMs remain available on the peer node.

  Upgrade-related alerts and error messaging have been enhanced to clearly explain restrictions and next steps.
- **Bug Fixes** – Bug fixes (see Bug Fixes).

New Features in Previous Releases

For a list of features introduced in:

- everRun Release 8.1.0.2 – See everRun Release 8.1.0.2 Release Notes.
- everRun Release 8.1.0.0 – See everRun Release 8.1.0.0 Release Notes.

Features No Longer Supported

Legacy Upgrade Paths No Longer Supported

Starting with everRun Release 8.1.1.0, upgrade methods that re-image a node are no longer supported. This includes the legacy migration path from CentOS to Ubuntu, and the legacy Ubuntu recovery re-image path.

Before attempting to upgrade, your system must be running at least the minimum supported everRun release of 8.0.0.0. If your system is on an older, CentOS-based release, you must first upgrade to a supported Ubuntu-based baseline (see Upgrading from everRun Release 7.x to 8.x or Higher). Then, apply the Release 8.1.1.0 upgrade kit from the everRun Availability Console.

Bug Fixes

Upgrade Now Verifies that Priv0 Is Isolated

When the private network (`priv0`) is not isolated, it can result in system management issues and degraded VM stability. However, in previous releases, the upgrade process proceeded when the private network (`priv0`) was not isolated.

In everRun Release 8.1.1.0, the upgrade process checks for external hosts on `priv0`. If `priv0` is not isolated, the upgrade stops and the following Alert is displayed: `The software upgrade failed for the following reason(s): Release information verification error.`

For information on resolving the `priv0` Network Isolation error before upgrading, see KB0017923 (*Release Information Validation Error - Resolving a priv0 Network Isolation Error During Upgrade Prequalification Checks*).

Restriction on Number of Virtual CPUs in Windows 11 Environments

In previous releases, in Windows 11 guest environments, the number of vCPUs could not exceed 8. Exceeding this limitation may have resulted in system hangs or data corruption.

This issue has been fixed in everRun.



Note: Before upgrading to Release 8.1.1.0, shut down all Windows 11 VMs so that the fix is correctly applied.

VM Creation Fails When Using a Remote ISO (NFS or CIFS) If the File Path Contains Spaces

Previously, the system used a validation check that did not allow whitespace characters in the remote ISO path. This caused VM creation to fail with an "Invalid file path" error if any folder or file name in the path included a space (for example, a folder named Windows IoT).

In everRun Release 8.1.1.0, spaces in remote ISO paths are now allowed, ensuring compatibility with standard naming conventions for network shares. Spaces in the `.iso` file name continue to be unsupported.

Upgrade Process Now Validates SSD Life Remaining

On everRun 200i and 250i systems running Release 8.1.0.x, the **Solid State Drive (SSD) Life remaining** check was not working correctly, and would report 100% life remaining when in fact more had been used.

To prevent outages and protect data, the upgrade process now detects and blocks upgrades when SSD life used is 95% or higher.

After upgrading to everRun Release 8.1.1.0, the **SSD Life remaining** reading in the everRun Availability Console is accurate. This appears on the **Sensors** tab of the **Physical Machines** page.

Cannot Attach USB Device to Unprotected VM

Previously, if you tried to attach a USB device to an unprotected VM on the **Virtual Machines** page of the everRun Availability Console, the operation failed.

This issue has been fixed in everRun Release 8.1.1.0.

Logical Disks Become Foreign on Node After Power Interruption

Previously, if both nodes in an everRun system lost power asynchronously, all logical disks in one of the nodes might become marked as foreign. The node recovery process preserved all data, but it re-created the `/boot` and `root` file systems and re-installed the everRun system software.

This issue has been fixed in everRun Release 8.1.1.0.

Bugs Fixed in everRun Release 8.1.0.2

For information, see Bug Fixes in everRun Release 8.1.0.2.

CVE Fixes

For a list of the CVE fixes, see Fixed CVEs.

Important Considerations

This section describes important information related to installing, upgrading, and running everRun Release 8.1.1.0:

- Upgrading to everRun Release 8.1.1.0
- Dell PowerEdge Systems and UEFI0074 Prompt During Installation
- Requirements for Upgrading Windows 10 Virtual Machine Guest Systems to Windows 11
- REST API Restrictions for Importing a VM
- SNMP Disabled by Default on everRun Systems
- Tested Guest Operating Systems

Upgrading to everRun Release 8.1.1.0

To upgrade to everRun Release 8.1.1.0, see Upgrading everRun Software Using an Upgrade Kit.

> **Caution**: Consider contacting your authorized Penguin Solutions service representative for assistance before upgrading your system from everRun Release 7.*x* to 8.*x*. They can help you to evaluate your system, ensure that it meets the prerequisites, locate any potential restrictions, and, if needed, discuss a strategy for the upgrade. If you encounter any problems during the upgrade process, contact your authorized Penguin Solutions service representative for assistance as soon as possible to avoid complications and potential data loss.

Dell PowerEdge Systems and UEFI0074 Prompt During Installation

During a fresh installation or autoinstall using the Ubuntu 24.04-based image on Dell PowerEdge servers, the system may display the following Dell UEFI firmware prompt during the second reboot of the everRun installation sequence:

```
UEFI0074: The Secure Boot policy has been modified since the last time the system was started. Review the Secure Boot policy settings in System Setup > System BIOS > Security Settings. If policy changes were expected, no action is necessary. Otherwise, modify the Secure Boot policy settings as needed.
```

This is an expected, informational Dell firmware notification, not a everRun installation failure or Secure Boot signature verification error. The notification occurs once during initial provisioning and does not reappear during normal system operation or subsequent reboots.

Dismiss the prompt and allow the system to proceed. Normal boot and installation continues without impacting BIOS security configurations. If the prompt halts an automated install pipeline, temporarily set Secure Boot to **Disabled** in Dell System Setup (**F2 > System BIOS > Security Settings**) during the initial deployment, and re-enable it after the OS installation completes.

Requirements for Upgrading Windows 10 Virtual Machine Guest Systems to Windows 11

Before you can upgrade a Windows 10 Virtual Machine guest system to Windows 11, the following requirements must be met on the Windows 10 guest system:

- BIOS Mode must be set to Unified Extensible Firmware Interface (UEFI).
- Secure Boot State and virtual Trusted Platform Module (vTPM) security must be enabled.

If the guest system does not meet these requirements, contact a Penguin Solutions service representative for assistance.

For information about enabling Secure Boot State and vTPM security, see Understanding Windows 11 Virtual Machine Differences from Other VMs.

REST API Restrictions for Importing a VM

If you import a virtual machine by using the **POST /v2/import/importvm** method of the REST API, observe the following restrictions:

- You can import a single OVF file only one time. If you try to import the same OVF file a second time, the import fails.
- The **importvm** method supports NFS, VCD, and USB import. It does not support PC or CIFS/SMB import.
- The REST API cannot validate input to the **importvm** methods. Input to the REST API is submitted as is. If the input is incorrect, the system returns an error message.
- If you import a Windows 11 VM, it must be the UEFI boot type, as described in Understanding Windows 11 Virtual Machine Differences from Other VMs.

SNMP Disabled by Default on everRun Systems

SNMP is disabled by default on everRun systems. Starting with everRun Release 7.9.0.0 or higher, the SNMP process is also stopped in the console operating system on each node. For security reasons, if you need to enable SNMP, you should disable SNMP v1 and v2, and enable only v3 by using the SNMP **Restricted** configuration. For details, see Configuring SNMP Settings and Security Hardening, as well as KB0015445.

Tested Guest Operating Systems

For a list of the guest operating systems tested with the current release, see Tested Guest Operating Systems. For information on guest operating systems tested or supported in previous releases, go to http://everrundoc.stratus.com, select the appropriate release, and then search for the guest operating system.

Known Issues

Fault-Tolerant VM Becomes Stuck After Rapid USB Detach and Re-Attach

Rarely, a VM can enter an inconsistent state if you detach a USB storage device from a VM, then re-attach it immediately. This occurs only when the VM is configured for fault tolerance, and the USB storage device is re-attached within one second (while synchronization is in progress and the standby node is becoming active).

To avoid this issue, do not attach the USB storage device immediately after detaching it from the VM. Wait until the VM finishes synchronizing.

If this issue occurs, follow the procedure below to restore the VM and re-enable fault tolerance:

1. Set the VM's preferred node to the node where the VM is currently active.
2. Power off the VM and wait until it is completely stopped.
3. Unplug the USB storage device from the node.
4. In the everRun Availability Console, wait until that disk no longer appears in the USB section on the **Physical Machines** page.
5. If automatic load balancing is required, reset the load balancing setting on the preferred node to **automatically balance** (see Load Balancing).
6. Power on the VM.
7. Wait until synchronization successfully completes and fault tolerance is restored.
8. If needed, plug the USB storage device into the node again.

This problem will be fixed in a future release.

Rolling Reboot Hangs Indefinitely When a VM Has an Attached USB Storage Device

Initiating a rolling reboot on a system where a VM has an attached USB storage device causes the rolling reboot process to stall indefinitely. The rolling reboot starts, but the process fails during node evacuation because VMs with attached USB storage devices cannot be migrated to the peer node.

To avoid this problem, before initiating a rolling reboot detach all USB storage devices from all VMs and verify that VM synchronization is healthy.

If this condition occurs, cancel or stop the rolling reboot, detach the USB storage device from the VM, wait for VM synchronization to complete, then restart the rolling reboot.

This problem will be fixed in a future release.

Backup of Stopped Virtual Machine Triggers Unnecessary Warning

When using the everRun Availability Console to back up a stopped VM, the following Warning appears: **This backup is crash consistent and not necessarily application consistent. Any inflight transaction will be available in subsequent backups. During the backup you may observe some performance impact.**

This warning is valid when backing up running VMs, but it is not applicable for stopped VMs, and you can safely ignore it. This issue will be fixed in a future release.

Duplicate VirtIO VCD Created After Power Cycles

Following power cycles of the system nodes, a duplicate VirtIO Virtual CD (VCD) may be created for a virtual machine. The extra VirtIO entry appears on the **Virtual CDs** page of the everRun Availability Console and remains stuck in the synchronizing state.

To resume normal operations, manually delete the duplicate VirtIO VCD from the VM configuration.

Intermittent Snaphot Creation Failure on Windows Server 2022 Virtual Machines

Creating a snapshot of a Windows Server 2022 guest VM may intermittently fail. This occurs when the guest operating system's Volume Shadow Copy Service (VSS) exceeds its 10-second write-hold threshold during snapshot preparation, particularly during periods of high disk I/O or shortly after container and disk resizing. This timeout is enforced internally by Windows VSS, so host-side timeout adjustments are not possible.

The workaround is to wait until guest I/O activity subsides and the VM is idle, then retry the snapshot operation.

Mismatch in OVA File Name Validation Between everRun Availability Console and REST API

Due to inconsistent string verification between the REST API and the everRun Availability Console, it is possible for Open Virtual Appliance (OVA) file names and VM names that contain special characters (such as **&**) to be imported via the REST API, even though they are restricted in the everRun Availability Console. These special characters can cause XML parsing errors that prevent system audit logs from displaying on the **Audit Logs** page of the everRun Availability Console.

To avoid this issue, ensure that OVA files and VM names adhere to standard everRun Availability Console naming restrictions, and omit special characters when invoking REST API import operations. See the naming conventions described in Creating a New Virtual Machine.

VM Import Succeeds With Invalid Destination Parameter

When using the REST API to import a virtual machine (VM) from a Virtual CD (VCD), the VM import operation succeeds even if invalid data is specified in the `destination` parameter. In a future release of everRun, the REST API will return an error if the `destination` parameter is invalid.

Swagger UI Accepts Invalid Authorization Tokens Without Immediate Error Message

When authorizing in the REST API Swagger UI (`/v2/login`), entering an invalid or improperly copied token allows authorization to proceed without displaying an error. The interface shows the session as authorized (displaying the **Logout** button), but subsequent API requests fail with a `401 Unauthorized` error. To avoid this issue, ensure that you copy the full, exact authentication token returned by `/v2/login` and paste it into the authorization field.

UEFI Debian VMs Fail to Boot After P2V Conversion

UEFI-based Debian virtual machines (versions 10.10, 11, and 12.8) may fail to boot following a Physical-to-Virtual (P2V) conversion. Although P2V conversion succeeds and the UEFI VM is created, the VM does not boot on the target device. Instead, boot processes stop in the UEFI shell.

P2V Conversion Fails for Rocky Linux Version 9.4

Physical-to-virtual migration (P2V conversion) is not supported for Rocky Linux 9.4 guests, because the version of **virt-v2v** (1.44.2) included in everRun Release 8.1.1.0 does not support that distribution. The process terminates with the error: `virt-v2v: error: virt-v2v is unable to convert this guest type (linux/rocky).`

A version of **virt-v2v** that supports Rocky Linux guests will be included in a future everRun release.

everRun Legacy BIOS Installation with Virtual Floppy Starts from `sdb` Disk

The everRun software installation typically configures the system to boot from the `sda` disk device. However, on a system that is configured with Legacy BIOS and a virtual floppy disk, `sda` is reserved for the floppy disk; therefore, the system boots from the `sdb` disk device.

Unable to Create Application-Consistent Snapshot for Ubuntu 24.04.1 VM

If you try to create an application-consistent snapshot of a VM with the Ubuntu 24.04.1 guest operating system, the system creates a crash-consistent snapshot even if you have installed the `qemu-guest-agent` package.

DVD/USB Recovery or Replacement Fails on System with 512e, 512n and 4K Logical Disks

If you use DVD/USB recovery (Recovering a Failed Physical Machine) or replacement (Replacing Physical Machines, Motherboards, NICs, or RAID Controllers) on a node in an everRun system with 512e, 512n, and 4K logical disks, the operation might fail with a `FATAL ERROR`. If this happens, contact your authorized Penguin Solutions service representative for assistance.

Guest Performance Issues with Large Guest Volumes

Guest volumes that are 2TB or greater in size and that become fragmented over time may significantly reduce guest performance. When volumes of this size are created, select **qcow2** as the disk image format. The format **qcow2**, instead of **raw**, reduces performance slightly, but prevents the significant performance reduction seen with **raw** guest volumes of this size. For information on creating guest volumes, see Creating a Volume in a Virtual Machine.

Graphical Installation Fails to Start When Installing Debian 11 Guest

If you install the Debian 11 guest operating system, the guest fails to start the graphic installation utility and the installation terminates with a fatal error. To resolve this issue, see KB0017035.

Removable Media and Migrating a PM or VM Using the P2V Client

Before migrating a PM or VM using the bootable P2V client (**virt-p2v**) ISO file, check if any removable media (for example, floppy disks, DVD drives, or external USB disks) are attached to the source image. If removable media are attached to the source image when you attempt to migrate a PM or VM, the error message **Conversion failed** appears. To prevent this issue, deselect the media in the **virt-p2v** window before starting the migration. To do so, access the **virt-p2v** window with the sections **Target properties** and **Fixed hard disks**, and then beneath **Fixed hard disks**, uncheck the box in the **Convert** column next to the removable media. See Migrating a Physical Machine or Virtual Machine to a System, particularly the section **To migrate a PM or VM to the everRun system**, for more information on using **virt-p2v**.

"The VM *name* has failed to start" Alert While Running the P2V Client Is Normal

While using the P2V client to migrate a VM from an everRun or ztC Edge system, it is normal if the source system displays the alert "The VM *name* has failed to start" during the migration process, because although the source VM is powered on and running the P2V client, the guest operating system does not start.

Unknown Error in everRun Availability Console When Importing an Unprotected VM

When you an import an unprotected VM (see Unprotected Operation), the everRun Availability Console might display an alert that "An unknown error has occurred." You can safely ignore this message and click **OK** to continue the import.

AVCLI Import or Restore of VM with Two VHDX Volumes Fails

If you use the **avcli vm-import** or **avcli vm-restore** commands to import or restore a virtual machine with two VHDX volume files, the operation fails. If needed, use the everRun Availability Console to import or restore the VM.

Maximum Resolution of a UEFI VM Console Session

On the **Virtual Machines** page of the everRun Availability Console, you can open a VM console session to display the console of the guest operating system running in the VM. When you open a console session to access a guest VM with a UEFI boot type, the console session may have a maximum resolution of 800x600. If needed, you can configure a higher resolution in the BIOS setup utility, as described in KB0017168.

Creating VCD Fails When Console Browser Is Microsoft Edge

When you are using Microsoft Edge as the browser for the everRun Availability Console, you cannot create a VCD: the process will fail. Instead, use another compatible browser (see Compatible Internet Browsers).

Mapping of Japanese Keyboards 106 and 109 for Console in Firefox May Be Incorrect

The mapping of the Japanese keyboards 106 and 109 may be incorrect when using Firefox to access the everRun Availability Console. Use Chrome or remote connection software (VNC or RDP), instead.

Cannot Enable SNMP Requests Without Traps

If you create an SNMP request in the everRun Availability Console, you must also create a trap; otherwise, the everRun Availability Console displays the error "Problem encountered updating SNMP. Make sure your settings are correct. Error: Configure SNMP failed." As a workaround, when creating a new SNMP request, click **Enable SNMP Requests** and **Enable SNMP traps**, do not define a Version 3 user, keep the default of **Restricted** requests, and specify at least one trap recipient. Adding a Version 3 user or clicking **Unrestricted** during the initial configuration might cause the configuration to fail. After the initial configuration is complete, you can then modify it to specify any settings that you require.

System Fails to Generate SNMPv3 Traps When Authentication Is Enabled

If you configure an SNMPv3 user and send SNMP traps with authentication enabled, the system fails to generate the SNMP traps.

VMs Running Windows 2016 with the Maximum vCPUs and Memory Will Not Reboot Cleanly

A Windows 2016 VM with the maximum supported number of vCPUs and the maximum amount of memory will not reboot cleanly. To avoid the problem, reboot the VM using the **Shutdown** button (on the **Virtual Machines** page, in the bottom pane for the VM), and then restart the VM using the **Start** button. To avoid the problem, reduce the number of vCPUs or the amount of memory assigned to the VM.

Some Browsers Unable to Connect a VNC When Using https

If you are connected to the everRun Availability Console using an **https** URL in some browsers, including the 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. In the Firefox browser, when the **Your connection is not secure** window appears:

1. Click **Advanced**. A message about an invalid security certificate appears.
2. Click **Add Exception**. The **Add Security Exception** dialog box appears with the console's location in **Location**.
3. Click **Confirm Security Exception**.

The VNC console appears.

Some Browsers Unable to Connect a VNC When Using http

If you are connected to the everRun Availability Console using an **http** URL in some browsers, including the Google® Chrome™ 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. If you cannot establish a connection, try using either the Mozilla Firefox or Microsoft Edge browser.

Reboot Required When Changing Node IP Address or Netmask Network Settings

When you change the IP address or netmask settings of a node as described in Configuring IP Settings, both the old and new settings are in effect until you reboot the node. Having both settings active may cause routing or connection issues.

Thermald “Unsupported CPU model or platform” Messages During Node Startup

On some systems, `syslog` may show `thermald` (thermal daemon) messages at node or system startup such as `Need Linux PowerCap sysfs` and `Unsupported cpu model or platform`. These messages are expected, and indicate that `thermald` is present but not needed for your existing hardware. The messages do not indicate a fault or thermal risk.

everRun Installation Fails on First Boot on a System with Virtual Drives

If you install the everRun software on a system with virtual drives created and initialized with `Fast Initialization`, the installation fails with a fatal error. For example:

```
FATAL ERROR: Error 5 creating everrun-root lv
```

To work around this issue, reboot the system and continue with the installation. The installation program will try to continue and perform one more reboot before the actual installation starts.

everRun Installation Fails on a Previously Installed System with EFI Boot Entries

If you install the everRun software on a system that was previously installed with everRun, the installation might fail with `grub-install` errors on the physical console of the system because it cannot overwrite the original EFI boot entries. For example, the `grub-install` error messages might appear similar to the following:

```
Stderr: Installing grub to /boot/efi.
Installing for x86_64-efi platform.
grub-install: warning: Cannot set EFI variable Boot000A.
grub-install: warning: efivarfs_set_variable: writing to fd 8
failed: Invalid argument.
grub-install: warning: _efi_set_variable_mode: ops->set_variable()
failed: Invalid argument.
grub-install: error: failed to register the EFI boot entry:
Invalid argument.
```

To work around this issue:

1. Reboot the system and open the BIOS configuration utility.
2. Remove the original EFI boot entries from the BIOS settings.

3. Save the updated settings and restart the system.
4. Restart the installation process.

For information about accessing and modifying BIOS configuration settings for a system, see the vendor documentation for the system hardware. If you need assistance, contact your authorized Penguin Solutions service representative.

Cannot Attach USB Device to VM on First Attempt

If you try to attach a USB device to a VM on the **Virtual Machines** page of the everRun Availability Console, the system might fail to detect valid partitions and attach the device on the first attempt. To work around the issue, disconnect and reinsert the USB device.

Incomplete Messages When Creating or Importing a Windows 11 Guest with the BIOS Boot Type

When you create or import a VM with the Windows 11 guest operating system, the VM must have the UEFI boot type as described in Understanding Windows 11 Virtual Machine Differences from Other VMs. If you create or import a Windows 11 guest with the BIOS boot type, the system appears to create the new VM successfully, but the VM cannot start. If you attempt to start the VM, the system automatically shuts down the VM and displays a message on the **Summary** tab of the **Virtual Machines** page to explain that the **VM *name* cannot start because it has a misconfigured boot type BIOS. To select the correct boot type (UEFI), you must delete the VM and create a new one**.

When this happens, the **Dashboard** and **Alert History** page might display incomplete information, such as a blank message, or a message that incorrectly states **VM *name* failed to boot and will be restarted** though the VM will be shut down. The messages also fail to explain how to create a VM with the correct boot type (UEFI).

Windows Server 2016 Guest Stuck in Booting State After Installation in Localized German Language

After installing Windows Server 2016 with the localized German version of the Windows Installer (ISO) in a VM with a single virtual network adapter, the VM is stuck in the **booting** state because its network is not working. To resolve this issue, see KB0017020.

Accessing Knowledge Base Articles

The **Penguin Solutions 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*nnnnnn*). You can access the **Penguin Solutions Customer Service Portal** and its Knowledge Base by using your existing portal credentials, or by creating a new user account, as follows.

To access the Knowledge Base

1. Log on to the **Penguin Solutions Customer Service Portal** at https://support.penguinsolutions.com.

   If needed, create a new account as follows:

   a. Click **Register**.

   b. Enter your contact information including your company email address and registration code, and then click **Submit**.

      Your company email address must include a domain name (for example, stratus.com) for a company that is a registered customer of Stratus. The portal sends an email to administrators of the company's account to approve the request.

   c. Upon approval, click the link in the email that you receive from Stratus.

   d. Enter a new password and finish configuring your account.

   If you need assistance creating an account, contact your authorized Penguin Solutions service representative.

2. In the portal, do one of the following:

   - In the **Search** box, enter keywords or the KB article number (KB*nnnnnn*) associated with the information you need, and then click the search button.
   - Click **Knowledge**, click the name of a product, and then browse available articles.

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.penguinsolutions.com/services-support/downloads/?tab=everrun. You can also search the **Knowledge Base** in the **Penguin Solutions Customer Service Portal** at https://support.penguinsolutions.com.

If you cannot resolve your questions with these online resources, and the system is covered by a service agreement, contact your authorized Penguin Solutions service representative. For information, see the **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 Development Kit (JDK) version 17.0.8.1 installed by typing:

  ```
  java -version
  ```

  If the client computer already has this version of JDK installed, the output looks similar to this:

  ```
  openjdk version 17.0.8.1 2023-08-24
  OpenJDK Runtime Environment (build 17.0.8.1+1-Ubuntu-
  0ubuntu122.04)
  OpenJDK 64-Bit Server VM (build 17.0.8.1+1-Ubuntu-
  0ubuntu122.04, mixed mode, sharing)
  ```

  If the output shows that the client computer has an older version of JDK installed, download and install the correct version from https://www.oracle.com/java/technologies/downloads/#java17.

- You need a valid username and password for the everRun system. AVCLI scripts embed the username/password, so use access control lists (ACLs) to safeguard your new credentials. AVCLI commands are encrypted using SSL.

Installing the Linux Client

To download the AVCLI client for Linux:

1. Download the Linux client:
   a. Go to the **Downloads** page at https://www.penguinsolutions.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:

a. Go to the **Downloads** page at https://www.penguinsolutions.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:

- 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::Simple Library`. Otherwise, this error appears:

```
unless ($error) {

     my $xs = XML::Simple->new();

     $stdout_hash = $xs->XMLin($stdout,forceArray=>0);

     return 0;

}

return 1;
```

Asynchronous Command Delay

Commands that invoke an action on the everRun system are called *asynchronous* because the command completes before the action completes. This allows for complex scripting.

If you want a command to complete inline before proceeding to the next command, create a simple script and use the `—wait` option. For example:

```
$ cli -x -H eagles node-workon --wait node0
```

In this example, `cli` does not complete until VMs and the management port are failed over from node0 to node1, and node0 is in maintenance mode. Without the `—wait` option, the command completes when acknowledged but before the resources are migrated.

Output Formatting

AVCLI can create user-friendly command output and program-friendly XML output.

User-Friendly Command Output

AVCLI output is formatted for easy readability. For example:

```
$ avcli -u admin -p admin -H avance -x node-info
node:
 -> name : node0
-> id : host:o14
-> state: running
-> sub-state : nil
-> standing-state : maintenance
-> mode : maintenance
-> primary : false
-> manufacturer : Dell
-> model : Dell PowerEdge 2950
-> maintenance-allowed : true
-> maintenance-guest-shutdown : false
-> cpus : 8
-> memory : 4,288,675,840
virtual machines:
node:
-> name : node1
```

```
-> id : host:o406
-> state : running
-> sub-state : nil
-> standing-state : warning
-> mode : normal
-> primary : true
-> manufacturer : Dell
-> model : Dell PowerEdge 2950
-> maintenance-allowed : true
-> maintenance-guest-shutdown : true
-> cpus : 8
-> memory : 4,288,675,840
virtual machines:
virtual machine:
-> name : eagles1
-> id : vm:o1836
```



Program-Friendly XML Output

You can create program-friendly XML output by using the `-x` or `--xml` global option. For example:

```
$ avcli -u admin -p admin -H localhost -x node-info
<?xml version="1.0" encoding="utf-8" standalone="no"?>
<avance>
<node>
<name>node1</name>
<id>host:o55</id>
```

<state>running</state>

<sub-state/>

<standing-state>normal</standing-state>

<mode>normal</mode>

<primary>false</primary>

<manufacturer>Intel Corporation</manufacturer>

<model>S5520UR</model>

<maintenance-allowed>true</maintenance-allowed>

<maintenance-guest-shutdown>false</maintenance-guest-shutdown>

<cpus>2</cpus>

<memory>25706889216</memory>

<virtual-machines/>

</node>

<node>

<name>node0</name>

<id>host:o23</id>

<state>running</state>

<sub-state/>

<standing-state>normal</standing-state>

<mode>normal</mode>

<primary>true</primary>

<manufacturer>Intel Corporation</manufacturer>

<model>S5520UR</model>

<maintenance-allowed>true</maintenance-allowed>

<maintenance-guest-shutdown>false</maintenance-guest-shutdown>

<cpus>2</cpus>

```
<memory>25706889216</memory>
<virtual-machines>
<virtual-machine>
<name>MyVM</name>
<id>vm:o6417</id>
</virtual-machine>
</virtual-machines>
</node>
</avance>
```



Note: The schema definition is maintained between releases.

If you do **not** specify `-x` or `--xml` and the command returns an error, a verbose message appears. For example:

```
$ avcli -H eagles -u admin -p admin 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.throwNonExi
stentResource(ResourceDisambiguateServiceProvider.java:56)
at
com.avance.yak.cli.ResourceDisambiguateServiceProvider.getResource
Id(ResourceDisambiguateServiceProvider.java:81)
at
com.avance.yak.cli.Command.findResourceId(Command.java:80)
at
com.avance.yak.cli.CommandWithUnparsedAmbiguousResourcesInvokeEach
.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:

```
$ avcli -H eagles -u admin -p admin 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.throwNonExi
stentResource(ResourceDisambiguateServiceProvider.java:56)
at
com.avance.yak.cli.ResourceDisambiguateServiceProvider.getResource
Id(ResourceDisambiguateServiceProvider.java:81)
at
com.avance.yak.cli.Command.findResourceId(Command.java:80)
at
com.avance.yak.cli.CommandWithUnparsedAmbiguousResourcesInvokeEach
.execute(CommandWithUnparsedAmbiguousResourcesInvokeEach.java:65)
at
com.avance.yak.cli.Command.execute(Command.java:194)
at
com.avance.yak.cli.CommandLine.execute(CommandLine.java:649)
```

```
    at
    com.avance.yak.cli.Program.main(Program.java:94)
```

AVCLI Command Descriptions

Click each heading to see the full list of AVCLI commands in that group.

> **Note**: The Examples section for each command assumes that you have set up a command shortcut as described in Executing a Command.

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

snmp-v3-add-agent-user

snmp-v3-add-trap-recipient

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-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-group-info-v2

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` *name*	The user with rights to join to the specified domain.
`--password` *password*	The password of the user with rights to join to the specified domain. If you do not give the password, you are automatically prompted for it.
`--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` *name*	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 everRun 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 `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 *HH*:*MM*:*SS,* 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/Dubai
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 [core | minimal | medium | large | full]
```

Description

The `diagnostic-create` command creates a new diagnostic of the specified type.

Options

`core`	Kernel and process core files.
`minimal`	About a day of data (approximately 2 to 10 MB).
`medium`	About a week of data (approximately 10 MB).
`large`	All data except core files (approximately 50 MB).
`full`	All data including core files (approximately 60 MB).

diagnostic-delete

Usage

```
avcli diagnostic-delete diagnostics...
```

Description

The `diagnostic-delete` command deletes the specified diagnostic files.

Options

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.

Options

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.

Options

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:o10
```

```
$ avcli diagnostic-fetch buggrab:o10 buggrab:o11 buggrab:o12
```

diagnostic-info

Usage

```
avcli diagnostic-info diagnostics...
```

Description

The `diagnostic-info` command displays information about all diagnostics or only about those specified.

Options

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 `dialin-info` command displays information about dial-in configuration.

disk-move-to-group

Usage

```
avcli disk-move-to-group disk... storage-group
```

Description

The `disk-move-to-group` command moves one or more logical disks to a storage group.

Options

disk	One or more disks to be moved.
storage-group	A storage group.

ealert-config

Usage

```
avcli ealert-config recipients...
```

Description

The `ealert-config` command configures e-Alert support in everRun systems.

Options

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 be sent to the recipient `admin@my-domain.com`:

```
$ avcli ealert-config admin@my-domain.com
```

ealert-disable

Usage

```
avcli ealert-disable
```

Description

The `ealert-disable` command disables e-Alert.

ealert-enable

Usage

```
avcli ealert-enable
```

Description

The `ealert-enable` command enables e-Alert.

ealert-info

Usage

```
avcli ealert-info
```

Description

The `ealert-info` command displays information about the e-Alert configuration.

help

Usage

```
avcli help [command] [-all]
```

Description

The `help` command provides help about a specific command or lists all AVCLI commands.

Options

`-all`	Display detailed information for all commands.

Examples

To display general command usage and a list of all commands for which `help` provides information:

```
$ avcli help
```

To display information about a specific command (`storage-info`, in this example):

```
$ avcli help storage-info
```

To display detailed information about all commands for which `help` provides information:

```
$ avcli help -all
```

image-container-info

Usage

```
image-container-info [image-container]
```

Description

The `image-container-info` command displays information about all image containers (also known as *volume containers*) or optionally, about only the specified image containers. Specifically, the command displays information about the portion of the image container that is available to the guest operating system.

Options

image-container	The name of the image container. If you do not provide this argument, the command displays information about all image containers.

Examples

```
$ avcli image-container-info

image-container:

-> name : root

-> id : imagecontainer:o58

-> hasFileSystem : false

-> isLocal : true

-> size : 21,479,030,784

-> size-used : 21,479,030,784

-> storage-group : none

image-container:

-> name : root

-> id : imagecontainer:o31

-> hasFileSystem : false
```

```
-> isLocal : true
-> size : 21,479,030,784
-> size-used : 21,479,030,784
-> storage-group : none
image-container:
-> name : swap
-> id : imagecontainer:o36
-> hasFileSystem : false
-> isLocal : true
-> size : 2,151,677,952
-> size-used : 2,151,677,952
-> storage-group : none
image-container:
-> name : swap
-> id : imagecontainer:o66
-> hasFileSystem : false
-> isLocal : true
-> size : 2,151,677,952
-> size-used : 2,151,677,952
-> storage-group : none
image-container:
-> name : shared.fs_image_container
-> id : imagecontainer:o77
-> hasFileSystem : false
-> isLocal : false
-> size : 1,073,741,824
```

```
-> size-used : 1,073,741,824
-> storage-group : none
image-container:
-> name : win7_ent_x86_32_sp1
-> id : imagecontainer:o1360
-> hasFileSystem : false
-> isLocal : false
-> size : 2,684,354,560
-> size-used : 2,684,354,560
storage-group:
-> name : Initial Storage Group
-> id : storagegroup:o21
image-container:
-> name : boot-chom1
-> id : imagecontainer:o1690
-> hasFileSystem : true
-> isLocal : false
-> size : 42,949,672,960
-> size-used : 37,787,627,192
storage-group:
-> name : Initial Storage Group
-> id : storagegroup:o21
```

image-container-resize

Usage

```
image-container-resize --new-size size image-container
```

Description

The `image-container-resize` command increases the size of the image container; specifically, the portion that is available to the guest operating system. (An *image container*, also known as a *volume container*, is a system-wide container that holds volumes and snapshots.) You might want to increase the image container's size if you need to take snapshots and the container lacks sufficient free space to do so.

Options

`--new-size` *size*	The new image-container size. By default, *size* 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-n.n.n.n_n-nnn.kit
```

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:o24
```

kit-info

Usage

```
avcli kit-info [kit_id...]
```

Description

The `kit-info` command displays information about all kits (the default) or only the specified kits.

Options

kit_id	One or more upgrade kits about which to display information. The value is the kit ID.

Examples

You can issue the `kit-info` command to obtain the *kit-id* value for 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:o24`:

```
[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 **Penguin Solutions Customer Service Portal** at https://support.penguinsolutions.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:o24
```

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:o24
```

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 `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` *name*	Local group name.
`--permissions` *permission-type*	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

groups	Local user groups.

Examples

```
$ avcli local-group-delete unprivileged_users
```

local-group-edit

Usage

```
avcli local-group-edit [--name] [--permissions] group-name-or-sid
```

Description

The `local-group-edit` command edits an existing local user group. You cannot edit default groups (`admin`, `platform_admin`, `read_only`). Only users with the role **Administrator** (the group `admin`) can execute this command. Users with the role **Platform Admin** (the group `platform_admin`) or **Read Only** (the group `read-only`) cannot execute this command.

Options

`--name` *name*	New local group name.
`--permissions` *permission-type*	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.

Options

groups	Local user groups.

local-user-add

Usage

```
avcli local-user-add --username name --realname name --email
address [--password password] [--new-password password] [--local-
groups groups] [--permissions permission-types]
```

Description

The `local-user-add` command adds a new local user to the everRun system. If the user's password is not given, the user is automatically prompted for it. The user is prompted twice to verify that the password was entered correctly. Only users with the role **Administrator** (the group `admin`) can execute this command. Users with the role **Platform Admin** (the group `platform_admin`) or **Read Only** (the group `read-only`) cannot execute this command.

Options

`--username` *name*	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 as `--password`.
`--realname` *name*	The user's real name.
`--email` *address*	The user's email address.
`--local-groups` *groups*	Local groups for the user to join, in the form of a comma-separated list.
`--permissions` *permission-types*	Local user permissions, in the form of a comma-separated list.

Examples

```
$ avcli local-user-add --username bsmith --realname "Bob Smith" --
email bsmith@example.com --password secret --local-groups admin
```

```
$ avcli local-user-add --username bsmith --realname "Bob Smith" --
email bsmith@example.com --local-groups users1,users2 --
permissions ADD_USER,UPDATE_USER
```

local-user-delete

Usage

```
avcli local-user-delete users...
```

Description

The `local-user-delete` command deletes the specified local users. Only users with the role **Administrator** (the group `admin`) can execute this command. Users with the role **Platform Admin** (the group `platform_admin`) or **Read Only** (the group `read-only`) cannot execute this command.

Options

users	One or more local users.

Examples

```
$ avcli local-user-delete afjord
```

```
$ avcli local-user-delete afjord bsmith tkirch
```

local-user-edit

Usage

```
avcli local-user-edit user [--username name] [--realname name] [--email address] [--password password] [--new-password password] [--local-groups groups] [--permissions permission-types] user-name-or-sid
```

Description

The `local-user-edit` command edits an existing user. If you do not give the `--password` option, the password is not changed. If you give the `--password` option, the command prompts the user twice to verify that the password was entered correctly. Only users with the role **Administrator** (the group `admin`) can execute this command. Users with the role **Platform Admin** (the group `platform_admin`) or **Read Only** (the group `read-only`) cannot execute this command.

Options

`--username` *name*	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` *name*	The user's real name.
`--email` *address*	The user's email address.
`--local-groups` *groups*	Local groups for the user to join, in the form of a comma-separated list.
`--permissions` *permission-types*	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.

Options

user	One or more users about which to display information.

localvm-clear-mtbf

Usage

```
avcli localvm-clear-mtbf
```

Description

The `localvm-clear-mtbf` command brings half of a VM back into service after it has been removed from service for failing too many times.

mail-server-config

Usage

```
avcli mail-server-config --host host [--ssl] [--tls] [--port port_
number] [--username user_name] [--password password] [--sender
sender_email_address]
```

Description

The `mail-server-config` command configures the mail server.

Options

`--host` *host*	The domain name or IPv4 address of the SMTP server.
`--ssl`	The system uses SSL encryption when communicating with the SMTP server. You cannot specify this option with `--tls`.
`--tls`	The system uses TLS encryption when communicating with the SMTP server. You cannot specify this option with `--ssl`.
`--port` *port_number*	The port number to use when connecting with the SMTP server.
`--sender` *sender_email_address*	The email address of the user who send the email.
`--username` *user_name*	The name for authentication on the *host*.
`--password` *password*	The .password to use with *user_name* for authentication on the *host*.

Examples

The examples below configure the SMTP server as `mail.my-domain.com`.

The following example configures the server:

```
$ avcli mail-server-config --host mail.my-domain.com
```

The following example configures the server using the protocol TLS and port 587 for communication and using the username `admin` and the password `secret` for authentication when sending email:

```
$ avcli mail-server-config --host mail.my-domain.com --tls --
port 587 --username admin --password secret --sender
sample@gmail.com
```

The following example configures the server using the protocol SSL for communication and using the username `admin` and the password `secret` for authentication when sending email:

```
$ avcli mail-server-config --host mail.my-domain.com --ssl --
username admin --password secret
```

The following example configures the server using the protocol SSL for communication and using the username `admin` and no password for authentication when sending email; because the command does not include a password, a prompt for a password will appear after you issue the command:

```
$ avcli mail-server-config --ssl --host mail.my-domain.com --
username admin
```

mail-server-disable

Usage

```
avcli mail-server-disable
```

Description

The `mail-server-disable` command disables the mail server.

mail-server-enable

Usage

```
avcli mail-server-enable
```

Description

The `mail-server-enable` command enables the mail server.

mail-server-info

Usage

```
avcli mail-server-info
```

Description

The `mail-server-info` command displays information about the mail server configuration.

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` *name*	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 `media-delete` command deletes the specified media.

Options

media	The media to be deleted.

media-eject

Usage

```
avcli media-eject [--cdrom name] [vm...]
```

Description

The `media-eject` command ejects media from the specified virtual machines.

Options

`--cdrom` *name*	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` *name*	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.

Options

media	The media about which to display information.

media-insert

Usage

```
avcli media-insert --iso [--cdrom] [vm...]
```

Description

The `media-insert` command allows you to insert media into the specified virtual machines.

Caution: When you insert a VCD into a running, fault-tolerant (FT) VM, it prevents the everRun software from migrating the VM to a different physical machine in the event of a failure. To restore fault-tolerant operation, unmount and eject the VCD as soon as you finish using it.

Options

`--iso` *name*	The ISO image to insert.
`--cdrom` *name*	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  [--force] 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.

Note: Changing the MTU of a business network that is being used as network0 or by running VMs may cause a temporary loss of connection to the system, so you must use the `--force` option. If you do not use the `--force` option with such networks, the following message appears:

> Changing the MTU of business networks may cause a temporary loss of connection to the system. If you still wish to do so, please use `---force` to override.

Options

--force	Forces the change of the MTU size. Specify this option when you want to change the MTU size of a business network whether or not it is being used by running VMs. If you do not specify this option, the MTU size cannot be changed.
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 --force network0 4000
```

```
$ avcli network-change-mtu --force network0 9000
```

network-change-role

Usage

```
avcli network-change-role networks... role
```

Description

The `network-change-role` command changes the role of the specified network to the specified role.

Options

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

networks	One or more networks.

Output

The following example shows the settings for four networks, including the default MTU value of 1500 for A-Links.

```
avcli network-info
shared network:
        -> name            : sync_2003
       -> id               : sharednetwork:o2334
        -> fault-tolerant  : ft
        -> role            : a-link
        -> bandwidth        : 10 Gb/s
       -> mtu              : 1500
shared network:
        -> name            : network0
       -> id               : sharednetwork:o64
        -> fault-tolerant  : ft
        -> role            : business
        -> bandwidth        : 1 Gb/s
       -> mtu              : 1500
```

```
shared network:
       -> name            : sync_2004
      -> id               : sharednetwork:o2333
       -> fault-tolerant  : ft
       -> role            : a-link
       -> bandwidth        : 10 Gb/s
      -> mtu              : 1500
shared network:
       -> name            : priv0
      -> id               : sharednetwork:o65
       -> fault-tolerant  : ft
       -> role            : private
       -> bandwidth        : 1 Gb/s
      -> mtu              : 1500
```

node-add

Usage

```
avcli node-add [--wait]
```

Description

The `node-add` command adds a PM to an everRun system.

Options

`--wait`
 	
`-w`
	Wait for the command to complete.

node-cancel

Usage

```
avcli node-cancel pm
```

Description

The `node-cancel` command cancels a PM that is being imaged.

Options

pm	The PM to cancel.

node-config-prp

Usage

```
avcli node-config-prp --nic1 adapter --nic2 adapter node
```

Description

The `node-config-prp` command configures a PRP adapter on the specified PM with two physical adapters.

You must run this command twice: once to configure the adapter on the first PM, and once to configure the adapter on the second PM.

Options

`--nic1` *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.

Options

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.

Options

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.

Options

`--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.

Options

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] [--finalize] 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. Use the `--finalize` option to enable the node (*pm*) to exit maintenance mode automatically after it restarts successfully.

Options

`--force` `-f`	Override the shutdown warning.
`--wait` `-w`	Wait for the command to complete.
`--finalize` `-F`	Takes the node out of maintenance mode.
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.

Options

`--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 `ova-info` command displays information about the specified OVA files.

Options

filename`.ova`	One or more OVA files.

ovf-info

Usage

```
avcli ovf-info filename.ovf...
```

Description

The `ovf-info` command displays information about the specified OVF files.

Options

filename`.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` *address*	The owner's email address.
`--name` *name*	The owner's name.
`--phone` *number*	The owner's phone number.

Examples

```
$ avcli owner-config --email "Bob Smith" --email
bsmith@example.org --phone 800-555-1234
$ avcli owner-config --phone 800-555-1234
```

owner-info

Usage

```
avcli owner-info
```

Description

The `owner-info` command displays information about the owner of the everRun system.

pm-clear-mtbf

Usage

```
avcli pm-clear-mtbf
```

Description

The `pm-clear-mtbf` command clears a PM's MTBF from the user interface.

proxy-config

Usage

```
avcli proxy-config --port name [--username name] [--password
password] host
```

Description

The `proxy-config` command configures the everRun system to use a proxy server. If you do not specify a user name, AVCLI assumes that authentication is not required to access the proxy server. If you specify a user name but not a password, you are prompted for the password.

Options

`--port` *number*	The port number.
`--username` *name*	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.my-
domain.com
```

```
$ avcli --port 8080 --username user proxy.my-domain.com
```

proxy-disable

Usage

```
avcli proxy-disable
```

Description

The `proxy-disable` command disables proxy.

proxy-enable

Usage

```
avcli proxy-enable
```

Description

The `proxy-enable` command enables proxy.

proxy-info

Usage

```
avcli proxy-info
```

Description

The `proxy-info` command displays information about the proxy configuration.

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:o36`). 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. (To add an SNMPv3 recipient, issue the snmp-v3-add-trap-recipient command.)

Options

`--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 specify one or more recipients.
`--community` *name*	The name of the SNMP community.
`--port` *number*	The port to use for SNMP. The default is 162.
`--recipients` *recipient ...*	The list of hosts to which to send traps using SNMP version 2c.
`--recipients-v1` *recipient-v1 ...*	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.

snmp-v3-add-agent-user

```
avcli snmp-v3-add-agent-user --username username --security-level
security_level [--authentication-type type] [--authentication-
pass-phrase pass_phrase] [--encryption-type type] [--encryption-
pass-phrase pass_phrase]
```

Description

The `snmp-v3-add-agent-user` command adds an SNMPv3 user (*username*) with read-only access to the everRun system. Other SNMPv3 servers can then send an SNMPv3 request (for example, `snmpwalk`) to this user in order to retrieve the values of objects listed in the management information base (MIB) files.

The system supports only one SNMPv3 user. If an SNMPv3 user already exists on the system and you issue this command, the command does not add the user, but displays an error message.

Create the SNMPv3 user on both nodes.

Options

`--username` *username*	The name of a user who has access to the SNMPv3 agent. *username* must be unique.
`--security-level` *security-level*	The user's security level. Valid values are: • `auth` for authentication and no privacy: Messages are authenticated, but not encrypted. `--authentication-type` and `--authentication-pass-phrase` are required. `--encryption-type` and `--encryption-pass-phrase` are optional. • `noauth` for no authentication and no privacy: No security is applied to messages; messages are not authenticated or encrypted. `--authentication-type`, `--authentication-pass-phrase`, `--`

	`encryption-type`, and `--encryption-pass-phrase` are optional. • `priv` for authentication and privacy: Messages are authenticated and encrypted. `--authentication-type`, `--authentication-pass-phrase`, `--encryption-type`, and `--encryption-pass-phrase` are required.
`--authentication-type` *type*	The user's type of authentication. Valid values are: • `MD5`: Configure the message digest algorithm (MD5) as the user's authentication type. • `SHA`: Configure the secure hash algorithm (SHA) as the user's authentication type.
`--authentication-pass-phrase` *pass_phrase*	The user's required pass phrase, which is used to generate the secret authentication key. The *pass_phrase* must be a minimum of eight characters.
`--encryption-type` *type*	The user's type of encryption. Valid values are: • `AES`: Configure the Advanced Encryption Standard (AES) as the user's encryption type. • `DES`: Configure the data encryption standard (DES) as the user's encryption type.
`--encryption-pass-phrase` *pass_phrase*	The user's required pass phrase, which is used to generate the secret encryption key. The *pass_phrase* must be a minimum of eight characters.

Examples

The following example adds the agent user `agentUser1` to the system. The SNMPv3 messages that `agentUser1` sends will be authenticated and encrypted.

```
$ avcli snmp-v3-add-agent-user --username agentUser1 --
security-level priv --authentication-type MD5 --
authentication-pass-phrase agentUser1AuthPassPhrase --
encryption-type AES --encryption-pass-phrase
agentUser1EncriptPassPhrase
```

The following example adds the agent user `agentUser2` to the system. The SNMPv3 messages that `agentUser2` sends will be authenticated, but not encrypted.

```
$ avcli snmp-v3-add-agent-user --username agentUser2 --
security-level auth --authentication-type SHA --
authentication-pass-phrase agentUser2AuthPassPhrase
```

The following example adds the agent user `agentUser3` to the system. The SNMPv3 messages that `agentUser3` sends will not be authenticated or encrypted.

```
$ avcli snmp-v3-add-agent-user --username agentUser3 --
security-level noauth
```

snmp-v3-add-trap-recipient

Usage

```
avcli snmp-v3-add-trap-recipient --recipient recipient --username
username --security-level security_level [--authentication-type
type] [--authentication-pass-phrase pass_phrase] [--encryption-
type type] [--encryption-pass-phrase pass_phrase]
```

Description

The `snmp-v3-add-trap-recipient` command adds a recipient server (*recipient*) and a trap user (*username*) to the `CallHomeInfo.xml` file on the everRun system. The everRun system can then send SNMPv3 traps to the trap user, when the user exists on the recipient server.

Options

`--recipient` *recipient*	The server that is the recipient of SNMPv3 traps. Specify a domain name or an IPv4 address.
`--username` *username*	The name of the trap user on the recipient server to whom the everRun system sends SNMPv3 traps.
`--security-level` *security-level*	The user's security level. Valid values are: • `auth` for authentication and no privacy: Messages are authenticated, but not encrypted. `--authentication-type` and `--authentication-pass-phrase` are required. `--encryption-type` and `--encryption-pass-phrase` are optional. • `noauth` for no authentication and no privacy: No security is applied to messages; messages are not authenticated or encrypted. `--authentication-type`, `--authentication-pass-phrase`, `--`

	`encryption-type`, and `--encryption-pass-phrase` are optional. • `priv` for authentication and privacy: Messages are authenticated and encrypted. `--authentication-type`, `--authentication-pass-phrase`, `--encryption-type`, and `--encryption-pass-phrase` are required.
`--authentication-type` *type*	The user's type of authentication. Valid values are: • `MD5`: Configure the message digest algorithm (MD5) as the user's authentication type. • `SHA`: Configure the secure hash algorithm (SHA) as the user's authentication type.
`--authentication-pass-phrase` *pass_phrase*	The user's required pass phrase, which is used to generate the secret authentication key. The *pass_phrase* must be a minimum of eight characters.
`--encryption-type` *type*	The user's type of encryption. Valid values are: • `AES`: Configure the Advanced Encryption Standard (AES) as the user's encryption type. • `DES`: Configure the data encryption standard (DES) as the user's encryption type.
`--encryption-pass-phrase` *pass_phrase*	The user's required pass phrase, which is used to generate the secret encryption key. The *pass_phrase* must be a minimum of eight characters.

Examples

The examples below add recipient servers and trap users to the `CallHomeInfo.xml` file on the everRun system.

The following example adds the recipient server `snmp1.my-domain.com` and the trap user `myTrapUser1`. The trap messages that the system sends to `myTrapUser1` will be authenticated and encrypted.

```
$ avcli snmp-v3-add-trap-recipient --recipient snmp1.my-
domain.com --username myTrapUser1 --security-level priv --
authentication-type MD5 --authentication-pass-phrase
trapUser1AuthPassPhrase --encryption-type AES --encryption-
pass-phrase trapUser1EncriptPassPhrase
```

The following example adds the recipient server `snmp2.my-domain.com` and the trap user `myTrapUser2`. The trap messages that the system sends to `myTrapUser2` will be authenticated, but not encrypted.

```
$ avcli snmp-v3-add-trap-recipient --recipient snmp2.my-
domain.com --username myTrapUser2 --security-level auth --
authentication-type MD5 --authentication-pass-phrase
trapUser2AuthPassPhrase
```

The following example adds the recipient server `snmp3.my-domain.com` and the trap user `myTrapUser3`. The trap messages that the system sends to `myTrapUser3` will not be authenticated or encrypted.

```
$ avcli snmp-v3-add-trap-recipient --recipient snmp3.my-
domain.com --username myTrapUser3 --security-level noauth
```

storage-group-create

Usage

```
avcli storage-group-create storage-group-name [--disk-type disk-type] [storage-group-name [--disk-type disk-type]]...
```

Description

The `storage-group-create` command creates one or more storage groups.

Options

storage-group-name	The name of the storage group to be created.
disk-type	Disk type of storage group to be created. Valid values are `512n` (default), `512e`, or `4k`.

storage-group-delete

Usage

```
avcli storage-group-delete storage-group-name...
```

Description

The `storage-group-delete` command deletes one or more storage groups.

Options

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.



Note: In everRun Release 7.8 or higher, the `storage-group-info` command is deprecated in favor of the `storage-group-info-v2` command, which displays more detailed information about 512n, 512e, and 4k storage groups. See storage-group-info-v2.

Options

`--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.
`--orphan`	Show the logical disks that are not part of any storage group.

Examples

```
$ avcli storage-group-info
storage-group:
     -> name : 512n_storageGroup
     -> description :
     -> id : storagegroup:o2945
     -> size : 0.00
     -> size-used : 0.00
     -> sector-size : 512 B
```

```
storage-group:
     -> name : 4k_storageGroup
     -> description :
     -> id : storagegroup:o3040
     -> size : 500.00 GiB
     -> size-used : 31.93 GiB
     -> sector-size : 4 KB
storage-group:
     -> name : Initial Storage Group
     -> description :
     -> id : storagegroup:o157
     -> size : 556.88 GiB
     -> size-used : 224.35 GiB
     -> sector-size : 512 B
storage-group:
     -> name : 512e_storageGroup
     -> description :
     -> id : storagegroup:o2976
     -> size : 556.88 GiB
     -> size-used : 0.00
     -> sector-size : 512 B
```

storage-group-info-v2

Usage

```
avcli storage-group-info-v2 [--disks] [--volumes] [storage-group...] [--orphan]
```

Description

The `storage-group-info-v2` command displays information about all storage groups, or optionally, only about those specified.

Note: In everRun Release 7.8 or higher, the `storage-group-info` command is deprecated in favor of the `storage-group-info-v2` command, which displays more detailed information about 512n, 512e, and 4k storage groups.

Options

`--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.
`--orphan`	Show the logical disks that are not part of any storage group.

Examples

```
$ avcli storage-group-info-v2
storage-group:
     -> name : 512n_storageGroup
     -> description :
     -> id : storagegroup:o2945
     -> size : 0.00
     -> size-used : 0.00
     -> logical-sector-size : 512 B
```

```
     -> physical-sector-size : 512 B
     -> disk-type : 512n
storage-group:
     -> name : 4k_storageGroup
     -> description :
     -> id : storagegroup:o3040
     -> size : 500.00 GiB
     -> size-used : 31.93 GiB
     -> logical-sector-size : 4 KB
     -> physical-sector-size : 4 KB
     -> disk-type : 4k
storage-group:
     -> name : Initial Storage Group
     -> description :
     -> id : storagegroup:o157
     -> size : 556.88 GiB
     -> size-used : 224.35 GiB
     -> logical-sector-size : 512 B
     -> physical-sector-size : 512 B
     -> disk-type : 512n
storage-group:
     -> name : 512e_storageGroup
     -> description :
     -> id : storagegroup:o2976
     -> size : 556.88 GiB
     -> size-used : 0.00
     -> logical-sector-size : 512 B
     -> physical-sector-size : 4 KB
     -> disk-type : 512e
```

storage-info

Usage

```
avcli storage-info [--disks] [--volumes] [storage-group...] [--
orphan]
```

Description

The `storage-info` command displays information about all storage groups, or optionally, only about those specified.

Options

`--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.
`--orphan`	Show the disks that are not part of any storage group.

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
```

Options

`true`	Enables VMs to automatically return to a node that has returned to service. .
`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`*.

Options

`--cluster-address` *`IP_address`*	The IP address of the everRun system.
`--static`	Explicitly sets the values, if specified, for `--prefix`, `--node0-address`, `--node0-gateway`, `--node1-address`, `--node1-gateway`, and `--dns-servers`; otherwise, DHCP sets these values (except `--cluster-address`).
`--prefix` *`prefix`*	The size of the network prefix. Values, in bits, are `8` (Class A), `16` (Class B), and `24` (Class-C).
`--node0-address` *`IP_address`*	The IP address of node0.
`--node0-gateway` *`IP_address`*	The IP address of the node0 gateway.
`--node1-address` *`IP_address`*	The IP address of node1.

`--node1-gateway` *IP_address*	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.

Options

`--deny`	Deny acceptance of the EULA.

unit-eula-reset

Usage

```
avcli unit-eula-reset
```

Description

The `unit-eula-reset` command resets the EULA acceptance state on an everRun system.

unit-info

Usage

```
avcli unit-info
```

Description

The `unit-info` command displays information about the specified everRun system.

unit-shutdown

Usage

```
avcli unit-shutdown
```

Description

The `unit-shutdown` command shuts down an everRun system.

unit-shutdown-cancel

Usage

```
avcli unit-shutdown-cancel
```

Description

The `unit-shutdown-cancel` command cancels a pending shutdown for an everRun system.

unit-shutdown-state

Usage

```
avcli unit-shutdown-state
```

Description

The `unit-shutdown-state` command returns the everRun system's shutdown state.

unit-synced

Usage

```
avcli unit-synced [--wait]
```

Description

The `unit-synced` command returns true if the everRun system is synchronized between all PMs; otherwise, it returns false.

Options

`--wait`
 	
`-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` *name_or_OID*	The name or OID of the VM.
`--deviceId` *device_Id*	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` *name*	The name or ID of a VM.
`--node` *node*	The name or ID of a PM to enable.

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` *name*	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` *priority*	The boot priority; values are 1 to 1000.
`--application-start-time` *minutes*	The estimated start time of the VM and application, in minutes. The minimum value is one minute.
`--autoStartMode` *autoStartMode*	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.

Examples

```
$ 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` *iso*	The ISO image to boot.
`--wait`	Wait for the VM to boot.
vm	One or more VMs that are being started.

Examples

```
$ avcli vm-cd-boot --iso MyISO vm1

$ avcli vm-cd-boot --iso MyISO vm:o100

$ avcli vm-cd-boot --iso MyISO --wait vm1
```

vm-copy

Usage

```
avcli vm-copy --source-vm source --name name [--description
"description"] [--cpu number] [--memory memory] [--availability
level] [--copy-volumes volumes] [--add-volumes volumes] [--keep-
volumes volumes] [--interfaces networks] [--storage-group group]
[--no-auto-start]
```

Description

The `vm-copy` command copies a VM from the specified VM. If any parameter is not specified, the corresponding value from the source VM is used.

Options

`--source-vm` *source*	The name or ID of the source VM.
`--name` *name*	The name of the VM to create.
`--description` "*description*"	The description for the new VM.
`--cpu` *number*	The number of virtual CPUs to assign to the VM.
`--memory` *memory*	The amount of memory, in megabytes, to assign to the VM.
`--availability` *level*	The availability level, high availability (`ha`) or fault tolerant (`ft`).
`--copy-volumes` *volumes*	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 *volume* 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.
`--add-volumes` *volumes*	The list of volumes to attach to this VM. A *volume* 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).
`--keep-volumes` *volumes*	The idle volumes to be attached to the new VM. Specify a volume either by name or ID.
`--interfaces` *networks*	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.
`--no-auto-start`	If set, the VM is not started after the copy has finished.

Examples

Copy a VM named `vm:o2046` to a new vm called `new_vm_name` and retain all of the original settings.

```
$ avcli vm-copy --source-vm vm:o2046 --name new_vm_name
```

Copy a VM named `vm_source` to a new high availability VM named vm_copy with 2 CPUs and 1,024 MB of memory. Copy the boot volume `volume:o7652` to the storage group `storagegroup:o129` with the new name `vm_source_vol0_bootable_copy`, image type of qcow2, and sector size of 512 B. Copy the volume `volume:o7749` with the default values. Also, create a new volume named `vm_copy_add_new1` with size 20GB to `storagegroup:o1090` and set the image type to qcow2 and sector size to 4 kB.

```
$ avcli vm-copy --source-vm vm_source --name vm_copy --cpu 2 --
memory 1024 --availability ha --copy-volumes
volume:o7652,storagegroup:o129,vm_source_vol0_bootable_
copy,qcow2,512 volume:o7749 --add-volumes
20GB,storagegroup:o1090,vm_copy_add_new1,qcow2,4096
```

Copy a VM named vm_source to a new high availability VM named new_vm_name with 2 CPUs and 1,024 MB of memory. Copy the boot volume boot_volume to the Initial-Storage-Group with the new name boot_volume_copy and the image type of raw. Copy the volume volume:o10158 with the default values. Create a new volume named volume_new1 with size 20GB to storagegroup:o71 and set the image type to qcow2 and sector size to 4 kB. Attach two idle volumes, volume_idle and volume:o19656. In addition, configure network interfaces network0 and sharednetwork:o61, set the default storage group to storagegroup:o71, 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-auto-start
```

vm-create

Usage

```
avcli vm-create --name name --cpu number --memory memory [--boot-type interface] --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.

Options

`--name` *name*	The name of the VM to create.
`--cpu` *number*	The number of virtual CPUs to assign to the VM.
`--memory` *memory*	The amount of memory, in megabytes, to assign to the VM.
`--boot-type` *interface*	The boot interface for the VM, either `bios` (the default) or `uefi`.
`--cdrom` *cd-name*	The CD-ROM from which you initially boot the VM. You cannot specify this option with `--kickstart` or `--remote-file-path`.
`--kickstart` *template*	The kickstart template to use when booting the VM. You cannot specify this option with `--cdrom` or `--remote-file-path`.
`--remote-file-path` *path*	A remote ISO repository to use when booting the VM. You cannot specify this option with `--cdrom` or `--`

	`kickstart`.
`--remote-type` *type*	The type of remote ISO repository specified in the `--remote-file-path` option. Valid options are `samba` or `nfs`.
`--remote-username` *username*	The user account to specify for access to the remote ISO repository specified in the `--remote-file-path` option. Required for samba repositories.
`--remote-password` *password*	The user password to specify for access to the remote ISO repository specified in the `--remote-file-path` option. Required for samba repositories.
`--availability` *level*	The availability level, high availability (`ha`, the default) or fault tolerant (`ft`).
`--interfaces` *networks, 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` *networks, MAC address*	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, specify 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` *volumes*	List of volumes to attach to this VM. A *volume* 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. ▪ 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, the BIOS boot interface, 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 --boot-
type bios \
--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, the UEFI boot interface, 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 --boot-
type uefi \
--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 `vm-create-from-snapshot` command creates a new VM from a VM snapshot.

Options

`--vm-snapshot-oid` *oid*	The vm snapshot OID from which to create the VM.
`--name` *name*	The name of the VM to create.
`--cpu` *number*	The number of virtual CPUs to assign to the VM.
`--memory` *memory*	The amount of memory, in megabytes, to assign to the VM.
`--availability` *level*	The availability level, high availability (`ha`) or fault tolerant (`ft`).
`--interfaces` *networks*	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` *volumes*	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` *prefix*	Adds the specified *prefix* 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 specify `--volume-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 centos-
boot 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.

Examples

```
$ 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] [--format format] [--volumes
volumes] [--wait] [--force] vm-name
```

Description

The `vm-export` command exports a VM in OVF/VHD or OVF/VHDX format to the directory specified by *pathname*. The command first exports VHD or VHDX 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

`--path` *pathname*	A pathname relative to the export mount point to where the exported OVF is written.
`--format` *format*	The format of the disk to be exported. Valid values are: • `vhd`–Virtual hard disk format. • `vhdx`–Hyper-V virtual hard disk format.
`--volumes` *name*	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.

Examples

```
$ avcli vm-export --path exports/excalibur1 excalibur1
```

```
$ avcli vm-export --volumes volume:o1345 volume:o1389 --path
exports/excalibur1 excalibur1
```

vm-import

Usage

```
avcli vm-import --archive filename.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.

> **Note**: You can use the `vm-import` command only to import OVF files that were exported from an everRun system. If you need to import a VMware OVF or OVA file, use the **Import/Restore Virtual Machine** wizard in the everRun Availability Console. For details, see Importing Virtualization Files.

Options

`--archive` *filename.ovf*	The OVF file archive to import.
`--no-auto-start`	Do not start the VM after the import has finished.
`--cpu` *number*	The number of CPUs to assign to the VM. This defaults to the value in the archive.
`--memory` *size*	The size of memory, in megabytes, to assign to the VM. This defaults to the value in the archive.
`--name` *vm-name*	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` *networks*	The list of shared networks to assign to the VM's interfaces. 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 the `--volumes` and `--storage-groups` rules are in effect.
`--volumes` *volumes*	Import only these volumes. By default, all available volumes from the OVF are imported.
`--volume-prefix` *prefix*	Adds the specified *prefix* 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 specify `--volume-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` *amount*	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` *level*	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 `qcow2` and `raw` (default).

Examples

```
$ avcli vm-import --archive vm1.ovf
$ avcli vm-import --archive vm1.ovf
$ avcli vm-import --name myVM --throttle low --archive vm1.ovf
$ avcli vm-import --cpu 2 --memory 1024 --archive vm1.ovf
$ avcli vm-import --interfaces network0 network1 --archive vm1.ovf
$ avcli vm-import --remap-volumes --archive vm1.ovf
$ avcli vm-import --storage-groups sm-0000 sm-0001 --archive
vm1.ovf
$ avcli vm-import --volumes boot_vol vol3 --data vol3 --archive
vm1.ovf
$ avcli vm-import --name myVM --protection-level HA --archive
vm1.ovf
$ avcli vm-import --archive vm1.ovf --image-format qcow2
```

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.

Examples

```
$ 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
networks
```

Description

The `vm-network-disable` command disables a VM's networks on a selected node.

Options

`--name` *name*	The name or ID of a VM.
`--node` *node*	The name or ID of a PM on which to disable the networks.
`--networks` *networks*	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 node1 \
--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` *name*	The name or ID of a VM.
`--node` *node*	The name or ID of a PM on which to enable the networks.
`--networks` *networks*	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 node1 \
--networks net2
```

vm-poweroff

Usage

```
avcli vm-poweroff [vm...] [--wait]
```

Description

The `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.

Examples

```
$ avcli vm-poweroff vm1

$ avcli vm-poweroff vm1 vm2

$ avcli vm-poweroff vm1 vm:o100
```

vm-poweron

Usage

```
avcli vm-poweron [vm...] [--wait]
```

Description

The `vm-poweron` command powers on the specified VMs.

Options

vm	One or more VMs to power on. Specify a VM either by name or ID.
`--wait` `-w`	Wait for the command to complete.

Examples

```
$ avcli vm-poweron vm1

$ avcli vm-poweron vm1 vm2

$ avcli vm-poweron vm1 vm:o100
```

vm-reprovision

Usage

```
avcli vm-reprovision --name name [--new-name name] [--description "description"] [--cpu number] [--cpu-emulation name] [--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.

Options

`--name` *name*	Specify the VM to reprovision. Reprovision only one VM at a time. Specify a VM either by name or ID.
`--new-name` *name*	Specify a new name for the VM.
`--description` "*description*"	Specify a description for the VM.
`--cpu` *number*	The number of virtual CPUs. This defaults to the VM's current amount.
`--cpu-emulation` *name*	The name of a supported emulation.
`--memory` *size*	The size of memory, in megabytes. This defaults to the VM's current amount.
`--addVolumes` *volumes*	The list of volumes to attach to this VM. A *volume* 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. ▪ The boot volume must be presented with 512 B as the sector size.
`--deleteVolumes` *volumes*	The list of volume(s) that are currently attached to the specified VM to be deleted. Specify a volume either by name or ID.
`--keepVolumes` *volumes*	The list of volume(s) that are currently attached to the specified VM to be kept attached to the VM. If you specify a volume that is currently attached but not specified in this list, the volume is detached (not destroyed) from the VM. Specify a volume either by name or ID.
`--interfaces` *networks, 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` *networks, MAC address*	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, specify the MAC address after the network name.
`--detach-boot-volume`	Detaches the VM's boot volume.
`--attach-boot-volume` *name*	Specify the name of a new boot volume for this VM. If the VM already has an attached boot volume, you must also specify `--detach-boot-volume`; otherwise, the command fails.

Examples

```
$ avcli vm-reprovision --name RHEL77 --cpu-emulation Haswell --cpu
1

$ avcli vm-reprovision --cpu 2 --name vm1

$ avcli vm-reprovision --cpu 2 --name vm:o100

$ avcli vm-reprovision --cpu 2 --memory 2048 --name vm:o100
```

Reprovision a VM named `vm001` that has a CPU, 512 MB of memory, a 1,024 MB volume and is attached to network0, and then allocate storage from `Pool-0001` for the volume. The volume is named `vm001_vol0`.

```
$ avcli vm-reprovision --cpu 1 --memory 512 --interfaces
network0 \

--addVolumes 1024,Pool-0001,vm001_vol0 --name vm1
```

Reprovision VM `vm1`, and then delete the volumes `volume:o411`, `data-vm1`, and `data-vm2` associated with it.

```
$ avcli vm-reprovision --deleteVolumes volume:o411 data-vm1
data-vm2 --name vm1
```

Reprovision VM `vm1` with the new data volume `data-1-7`, delete volume `volume:o1043`, keep volumes `volume:o1`, `volume:o2`, `volume:o4`, and attach network interfaces `sharednetwork:o129` and `sharednetwork:o130`.

```
$ avcli vm-reprovision --cpu 3 --memory 3359 --addVolume
2500,storagegroup:o54,data-1-7 --deleteVolumes volume:o1043 -
-keepVolumes volume:o1 volume:o2 volume:o4 --interfaces
sharednetwork:o129 sharednetwork:o130 --name vm1
```

Reprovision VM `vm1` with the same parameters from the previous example. Also rename the VM `vm2` and add a description.

```
$ avcli vm-reprovision --cpu 3 --memory 3359 --addVolumes
2500,storagegroup:o54,data-1-7,qcow2 --deleteVolumes
volume:o1043 --keepVolumes volume:o1 volume:o2 volume:o4 --
interfaces sharednetwork:o129 sharednetwork:o130 --name vm1 -
-new-name vm2 --description "This is the vm description"
```

Reprovision VM `vm001` with two CPUs, 2048 MB of memory, one new data volume `vm001_data1` in qcow2 format with 4 kB sector size, and keep volume `o7517`.

```
$ avcli vm-reprovision --cpu 2 --memory 2048 --addVolumes
20GB,storagegroup:o1090,vm001_data1,qcow2,4096 --keepVolumes
volume:o7517 --name vm001
```

Swap two VMs' boot disks.

Detach boot volume:

```
$ avcli vm-reprovision --detach-boot-volume --name p56xen4
```

Switch boot volume:

```
$ avcli vm-reprovision --detach-boot-volume --attach-boot-
volume boot-p56xen4 --name p56xen8
```

Attach a detached boot volume to a different VM:

```
$ avcli vm-reprovision --attach-boot-volume boot-p56xen8 --
name p56xen4
```

vm-restore

Usage

```
avcli vm-restore --archive filename.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 `vm-restore` command restores a VM from an OVF file.

Options

`--archive` *filename*`.ovf`	The OVF file archive to restore.
`--no-auto-start`	Do not start the VM after the restore has finished.
`--cpu` *number*	The number of CPUs to assign to the VM. This defaults to the value in the archive.
`--memory` *size*	The size of memory, in megabytes, to assign to the VM. This defaults to the value in the archive.
`--name` *vm-name*	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` *networks*	The list of shared networks to assign to the VM's interfaces. By default, values in the archive or available shared networks are assigned.
`--volume-prefix` *prefix*	Adds the specified *prefix* 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 specify `--volume-prefix new`, the imported volume would be `new-ocean_boot`.
`--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 restoring 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.

Examples

```
$ 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 `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.

Examples

```
$ avcli vm-shutdown vm1

$ avcli vm-shutdown vm1 vm2

$ avcli vm-shutdown vm1 vm:o100
```

vm-snapshot-create

Usage

```
avcli vm-snapshot-create [--volumes | --no-data][--description
"description"] [--desire] [--require] vm-name
```

Description

The `vm-snapshot-create` command creates a VM snapshot.

Two snapshot consistency levels are supported:

- *Crash consistency*: The restored data is in the same state that it would have been if the system had crashed at the exact moment that the snapshot was taken. A crash-consistent snapshot does not capture the contents of memory or any pending I/O operations.
- *Application consistency*: Before the snapshot is taken, cooperating applications are briefly frozen so that transactions are complete, buffers flushed, files closed, and so on. This ensures that cooperating applications start from a consistent state. This is the highest level of consistency.

Options

`--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 specify `--no-data`. If you specify `--no-data`, no volumes are included in the snapshot. These arguments are mutually exclusive.
`--description "`*description*`"`	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-create-
disable
```

vm-snapshot-create-enable

Usage

```
avcli vm-snapshot-create-enable
```

Description

The `vm-snapshot-create-enable` command enables the system's ability to create snapshots. By default, the system's ability to create snapshots is enabled. Only users with the role **Administrator** (the group `admin`) can execute this command. Users with the role **Platform Admin** (the group `platform_admin`) or **Read Only** (the group `read-only`) cannot execute this command.

Examples

```
$ avcli -H localhost -u admin -p password vm-snapshot-create-
enable
```

vm-snapshot-delete

Usage

```
avcli vm-snapshot-delete snapshot...
```

Description

The `vm-snapshot-delete` command deletes the specified snapshots.

Options

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 [--format format] snapshot
```

Description

The `vm-snapshot-export` command exports a snapshot of a VM in OVF/VHD or OVF/VHDX format to the directory specified by *pathname*.The command first exports VHD or VHDX 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

`--wait`	Wait for the export operation to complete. Specify this option to view the export progress.
`--silent`	Suppress progress output.
`--volumes` *volumes*	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 exported OVF is written.
`--format` *format*	The format of the snapshot to be exported. Valid values are: • `vhd`–Virtual hard disk format. • `vhdx`–Hyper-V virtual hard disk format.
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

vm	One or more VMs to unlock. Specify a VM either by name or ID.

Examples

```
$ avcli vm-unlock vm1
$ avcli vm-unlock vm:o100
```

vm-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` *name*	The name or ID of a VM.
`--node` *node*	The name or ID of a PM on which to disable the volumes.
`--volumes` *volumes*	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` *name*	The name or ID of a VM.
`--node` *node*	The name or ID of a PM on which to enable the volumes.
`--volumes` *volumes*	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 node1 \
--volumes volume:o2249
```

volume-info

Usage

```
avcli volume-info [volume...]
```

Description

The `volume-info` command displays information about all volumes, or optionally, only about those specified.

Options

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` *size*	The new volume size. By default, *size* 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

- Tested Guest Operating Systems
- Physical Machine System Requirements
- Important Physical Machine and Virtual Machine Considerations
- Accessing Knowledge Base Articles
- Creating a SplitSite Configuration
- Fixed CVEs
- REST API

Tested Guest Operating Systems

The following table lists the guest operating systems for virtual machines (VMs) that Stratus has tested on the current release of everRun software.

Notes:

- For the latest information about guest operating systems that have been tested with the everRun software, see the everRun Supported Guest Operating Systems web site at https://www.stratus.com/services-support/customer-support/software-support/everrun-supported-guest-operating-systems/.
- If you are upgrading from a previous release of everRun and your system is already running guest operating systems that were supported in the previous release, the guests will also be supported when you upgrade to the current release, even if they are not listed in the following table.
- For information about importing or migrating guest operating systems, see Importing Virtualization Files or Migrating a Physical Machine or Virtual Machine to a System.
- You can import or migrate a guest operating system from VMware ESXi™ only if the host is running VMware vSphere® Release 6.7, 7.0, or 8.0.

Operating System	Version(s)	Boot Firmware Interface
AlmaLinux OS 9	9.2, 9.3, 9.4 (all 64-bit)	BIOS UEFI
AlmaLinux OS 8	8.8 (64-bit)	BIOS UEFI
CentOS Stream 9	9.0 (64-bit)	BIOS UEFI
Debian 12	12.8 (64-bit)	BIOS UEFI
Debian 11	11.6, 11.11 (both 64-bit)	BIOS UEFI

Operating System	Version(s)	Boot Firmware Interface
Debian 10	10.9, 10.10 (both 64-bit)	BIOS UEFI
Microsoft Windows Server 2025 (Standard, Datacenter)	64-bit	BIOS UEFI
Microsoft Windows Server 2022 (Standard, Datacenter)	64-bit	BIOS UEFI
Microsoft Windows Server 2019 (Essentials, Standard, Datacenter)	64-bit	BIOS UEFI
Microsoft Windows Server 2016 (Essentials, Standard, Datacenter)	64-bit	BIOS UEFI
Microsoft Windows 11 (Home, Pro, Enterprise)	64-bit	UEFI [1]
Microsoft Windows 11 IoT Enterprise	64-bit	BIOS UEFI
Microsoft Windows 10 (Home, Pro, Enterprise)	64-bit	BIOS UEFI
Microsoft Windows 10 IoT Enterprise	64-bit	BIOS UEFI

[1]You must select the UEFI boot interface for the Windows 11 guest operating system to support its strict security requirements. For more information, see Understanding Windows 11 Virtual Machine Differences from Other VMs.

Operating System	Version(s)	Boot Firmware Interface
Oracle Linux 9	9.2, 9.4 (both 64-bit)	BIOS UEFI
Oracle Linux 8	8.7, 8.8, 8.9, 8.10 (all 64-bit)	BIOS UEFI
Red Hat Enterprise Linux 10 (Workstation, Server)	10.0, 10.1, 10.2 (all 64-bit)	BIOS UEFI
Red Hat Enterprise Linux 9 (Workstation, Server)	9.2, 9.3, 9.4, 9.6, 9.8 (all 64-bit)	BIOS UEFI
Red Hat Enterprise Linux 8 (Workstation, Server)	8.7, 8.8, 8.10 (all 64-bit)	BIOS UEFI
Rocky Linux 9	9.2, 9.4 (both 64-bit)	BIOS UEFI
Rocky Linux 8	8.8 (64-bit)	BIOS UEFI
SUSE Linux Enterprise Server (SLES)	12 SP2, 15 SP3, 15 SP5, 15 SP6 (all 64-bit)	BIOS UEFI
Ubuntu Server	20.04 LTS, 22.04 LTS, 24.04 LTS (all 64-bit)	BIOS UEFI

Physical Machine System Requirements

The following table presents the minimum and maximum capacities of the listed devices for physical machines running on everRun systems.

Notes:

- For the latest information about physical machines that have been tested with the everRun software, see the everRun Supported Hardware web site at https://www.stratus.com/services-support/customer-support/everrun-supported-hardware/.
- When qualifying a new processor family (for example, Emerald Rapids) with the everRun software, Stratus typically tests the software with at least two or three processor models from the same processor family to ensure compatibility. Although the following table lists only a few of the extensive list of processors from each family, if a processor family is listed here, the everRun software supports other processors from the same processor family.

Physical Device	Minimum	Maximum Tested	Notes
Processors/CPUs: Intel® Xeon® 6 P-Core (Granite Rapids) Intel® Xeon® Gold 63*XX*, Gold 53*XX* (Ice Lake) Intel Xeon Gold 62*XX*R, Gold 52*XX*R (Cascade Lake Refresh) Intel Xeon Gold 62*XX*, Gold 52*XX* (Cascade Lake) Intel Xeon Gold 61*XX*, 51*XX* (Skylake) Intel Xeon Gold 55*XX* (Emerald Rapids) Intel Xeon Gold 54*XX*S (Sapphire Rapids) Intel Xeon Silver 44*XX*Y	1	2	

Physical Device	Minimum	Maximum Tested	Notes
(Sapphire Rapids) Intel Xeon Silver 43*XX* (Ice Lake) Intel Xeon Silver 42*XX*R (Cascade Lake Refresh) Intel Xeon Silver 42*XX* (Cascade Lake) Intel Xeon Silver 41*XX* (Skylake) Intel Xeon Bronze 31*XX* (Skylake) Intel Xeon E5-2*XXX* v4 (Broadwell) Intel Xeon E5-2*XXX* v3 (Haswell) Intel Xeon E5-2*XXX* v2 (Ivy Bridge) Intel Xeon E5-2*XXX* (Sandy Bridge) Intel Xeon E5-1*XXX* v4 (Broadwell) Intel Xeon E5-1*XXX* v3 (Haswell) Intel Xeon E5-1*XXX* v2 (Ivy Bridge) Intel Xeon E5-1*XXX* (Sandy Bridge) Intel Xeon E3-1*XXX* v6 (Kaby			

Physical Device	Minimum	Maximum Tested	Notes
Lake) Intel Xeon E3-1*XXX* v5 (Skylake) Intel Xeon E3-1*XXX* v4 (Broadwell) Intel Xeon E3-1*XXX* v3 (Haswell) Intel Xeon E3-1*XXX* v2 (Ivy Bridge) Intel Xeon E3-1*XXX* (Sandy Bridge) Intel Xeon E-2*XXX* (Coffee Lake) Intel Xeon W-1290 TE (Comet Lake) Intel Xeon W-1250 TE (Comet Lake)			
Number CPU Sockets per Physical Machine	1	2	
Physical Memory	8 GB	384 GB	
Internal Disk Count per Physical Machine	2	24	Minimum 2 drives per PM for FT. VM disks/volumes are replicated on both PMs.
Total Disk Capacity	36 GB	9.4 TB	

Physical Device	Minimum	Maximum Tested	Notes
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 Management link.
Quorum Servers	0	2	

Important Physical Machine and Virtual Machine Considerations

For optimal implementation of physical machines and virtual machines, be aware of the configuration maximums and requirements described in the following sections:

- Physical Machine System Requirements
- 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 *n* vCPUs	*n* + 2 (typical)
Each HA guest with *n* vCPUs	*n* + 1 (typical)

Note: A non-hyperthreaded physical CPU core can handle 1 physical thread. A hyper-threaded physical CPU core can handle 2 physical threads.

The actual number of required threads depends upon the workload. The guidelines above should cover most workloads. Since any given workload may require fewer or more threads, it's good practice to test and characterize your specific workload.

Examples

A single 4-vCPU FT guest typically requires:

- 2 threads for host/system management
- 6 threads for guest
    - **8 threads total** (a single socket 4-core hyper-threaded system)

Four 5-vCPU FT guests typically require:

- 2 threads for host/system management
- 7 threads for first guest
- 7 threads for second guest
- 7 threads for third guest
- 7 threads for fourth guest
    - **30 threads total** (a dual socket 8-core hyper-threaded system)

Virtual Machine Limits

For systems with many or large virtual machines (VMs), configure everRun with 10 Gb sync links, and for the everRun software itself, 4 vCPUs and 4096 MB. Refer to the **Preferences -> Systems Resources** page

in the everRun Availability Console for instructions on setting the everRun system resources to the maximum.

The following table lists everRun system VM limits.

Item	Limits
Maximum vCPUs per FT VM	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	8
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
everRun System Disk	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 files (`-f 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. Use network-attach tape drives instead.
Console Connectivity	Each PM's text console is available for the host operating system. However, VGA mode is not supported; that is, the PM must be at run-level 3 and cannot run at run-level 5. See "System Management" below.
SSD Support	everRun supports solid state drives according to the storage controller vendor specifications.
System Management	everRun 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

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.

It is strongly recommended 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.

Note: This topic and its subtopics describe a SplitSite configuration with a quorum server. It is very important 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* (KB0014558), and also contact your authorized Penguin Solutions 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 *active* 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 *physical machines* (PMs) or *nodes*.) A-Links must be connected point-to-point, and A-Link traffic cannot be routed. A everRun system requires two A-Links. On some systems, these connections have blue and yellow cables (and ports). You can use VLAN connections for A-Links in a distributed local site 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 synchronized between the active node and the standby node. Each VM has its own AX pair (see VM, Active node, and Standby node)
Business network (ibiz)	A network connection from the 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 active simultaneously, which produces divergent copies of data within each active guest (see AX and VM). Split-brain can occur when all communication paths between node0 and node1 are disconnected (see Node0 and node1). Using the quorum service prevents a split-brain condition (see Quorum server).
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 tolerance, at the expense of more network setup and more extensive configuration 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 responsible for maintaining the overall state of the system. Determining which node is primary is part of system management (see Primary node). System management is also responsible for displaying information within the everRun Availability Console.

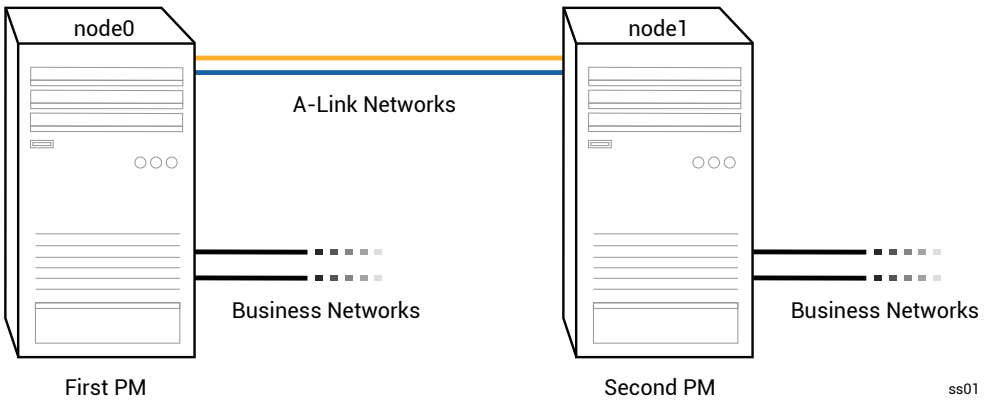
UPS	Uninterruptable power supply. An external battery backup for electrical equipment that prevents short power outages from affecting availability.
VLAN	Virtual LAN. A VLAN is a set of devices on one or more LANs that are configured to communicate as if they were attached to the same cabled network, when in fact they are located on different LAN segments. VLANs are configured at the network infrastructure level, not within the everRun system. In a SplitSite configuration, the A-Link connections are implemented as isolated VLANs (see A-Link).
VM	Virtual Machine (also referred to as a guest). A system typically has one or more VMs (or guests) allocated and running applications via guest operating systems.

Creating the Configuration

To create 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

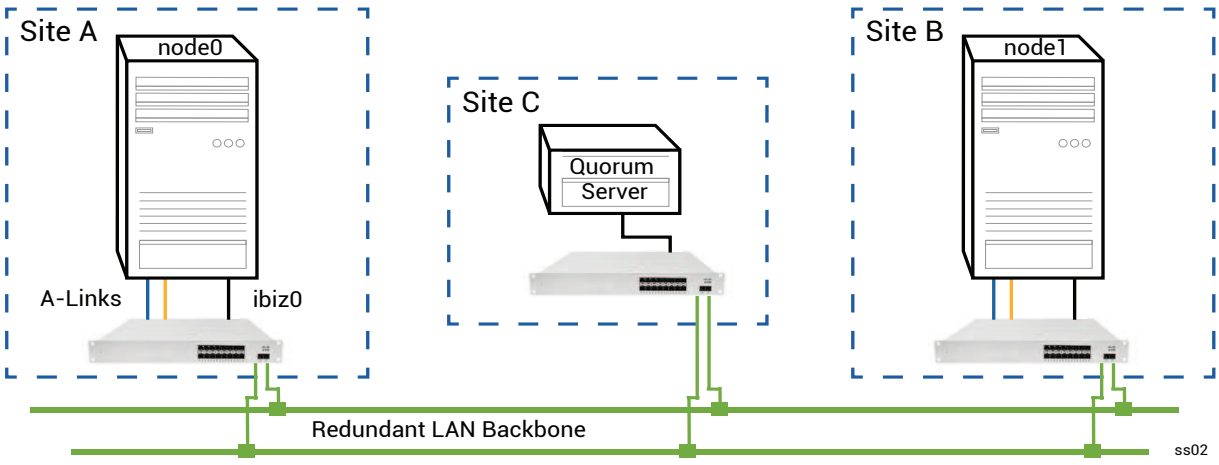
In a typical everRun system configuration, two PMs are directly connected by at least two network cables for A-Links. One A-Link typically serves as the private network (priv0). The two PMs have additional network connections for business networks, which the everRun Availability Console and guest VMs hosted by the system use. The following figure illustrates a typical configuration.



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.



Notes:

1. Each A-Link should be connected on its own VLAN configured between switch A and switch B.
2. DNS servers and gateways are not included in the illustrations, for clarity, but you must ensure that the SplitSite configuration includes a connection to a DNS server and a gateway in the event of a network failure.
3. For maximum protection, you should install redundant switches at each at site though the illustration does not show these switches. For the illustrated configuration, site A and site B would *each* include two switches. The A-Links are routed through one switch and the business networks are routed through the other switch. If possible, use different circuits to power the switches or use a UPS to prevent brief power loss failures.

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 a 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 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).
9. (Re-)join the VMs—Migrate the VMs from node to node (see Migrating a Physical Machine or Virtual Machine to a System). Verify the correct network failover of VM networking.
10. Assess the status of network and validate Ethernet failover (see the Networks Page).

Meeting Network Requirements

This topic describes the network requirements and considerations of A-Links, business networks, 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.
- The two nodes can access the quorum server using IPv4 network addressing.

Private network connections (priv0 through priv*n*) 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.If possible, configure two quorum servers. 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 service is lost *before* a node is lost, and if both nodes can continue to contact the alternate quorum server, the selection is moved to the alternate quorum server. Fault handling occurs in a context of the selected quorum server only.

If you create an alternate quorum service, you need to add a second quorum IP address when adding the quorum service in the 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 2019, Windows Server 2016, Windows Server 2012, 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.penguinsolutions.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.

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\QS
ServerPortForReceive
```

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 (recommended).
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**.
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.

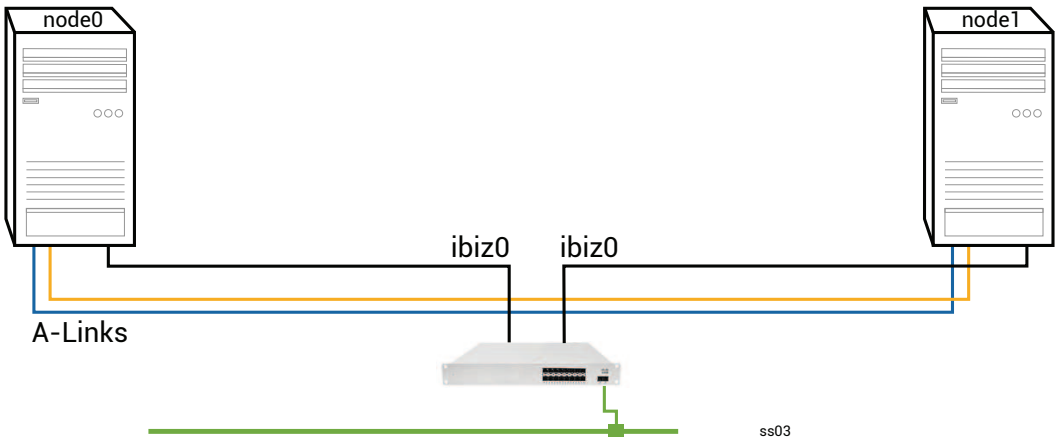
An everRun system is designed to provide high availability for one or more guest VMs, which allows the VMs to continue to run even during failures that would otherwise create application downtime. The everRun system can continue to run guest VMs even with, for example, the loss of a single network connection, 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.

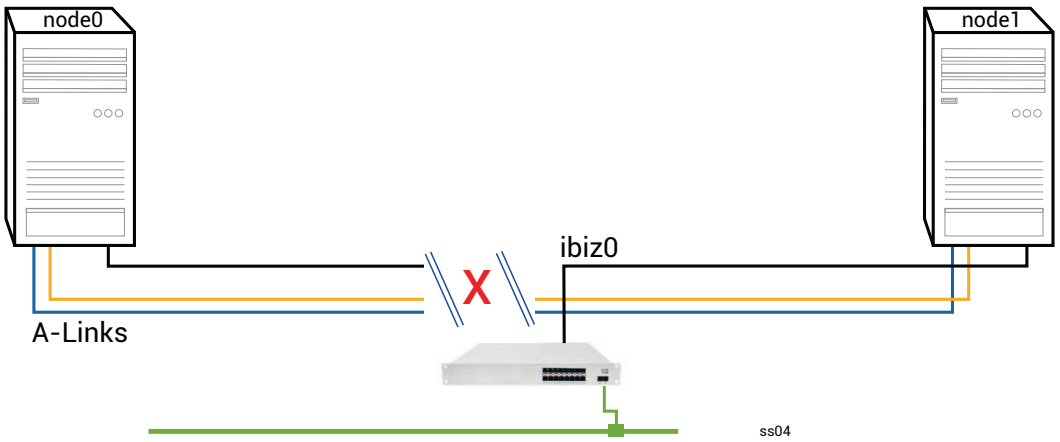
Example 1: A System Without a Quorum Server Experiences a Split-brain Condition

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 (faultless) operation. The following illustration shows normal connections.



A Catastrophic Fault

A careless fork-truck operator crashes through the wall, severing all of the network connections (both business and A-Links), while leaving the power available and the system running. The following illustration shows the fault condition.



Fault Handling

The two nodes handle the fault, as follows:

- Node0—The AX on node0 detects the loss of both A-Links as well as all other network paths. Since the node0 AX can no longer detect the presence of its partner, the node0 AX becomes active and runs the guest VM. The application inside the guest VM continues to run, perhaps in a limited capacity due to the loss of the network.
- Node1—The AX on node1 also detects the loss of both A-Links, but ibiz0 remains available. As its partner does not respond to messages on ibiz0, the node1 AX is now active. The application inside the guest VM continues to run, perhaps not noticing any problems with the system.

From the perspective of an application client or an external observer, the guest VMs are both active and generate network messages with the same return address. Both guest VMs generate data and see different amounts of communication faults. The states of the guest VMs becomes more divergent over time.

Recovery and Repair

After some time, network connectivity is restored: the wall is repaired and the network cables are replaced.

When each AX of the AX pair realizes that its partner is back online, the AX pair with the fault handler rules choose the AX that continues as active. The choice is unpredictable and does not include any consideration for which node's performance was more accurate during the split-brain condition.

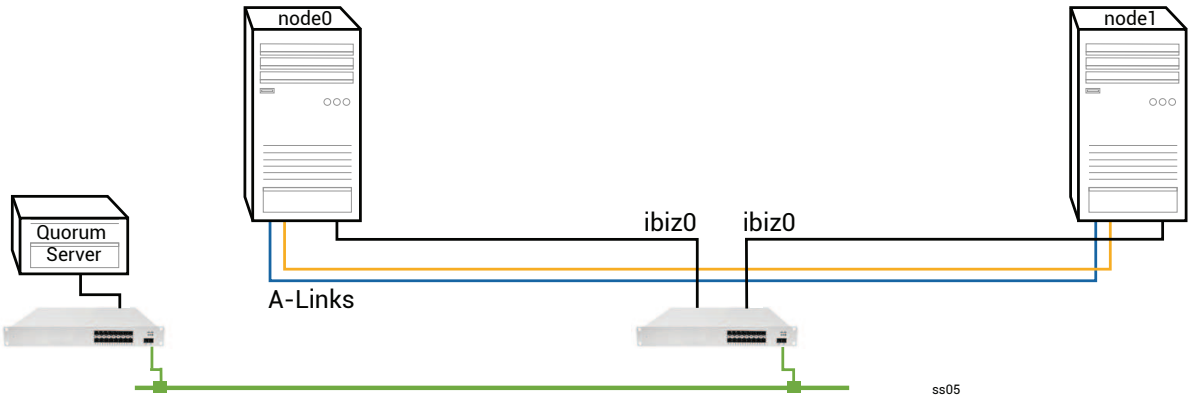
The data that was generated from the (now) Standby node is overwritten by the resynchronization of the Active node, and thus the data on the (now) Standby node is lost forever.

After a split-brain condition, the system requires several minutes to resynchronize, depending on how much disk activity needs to be sent to the standby node. If several guest VMs are running with different Active nodes, synchronization traffic may occur in both directions.

Note: In some cases, the 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

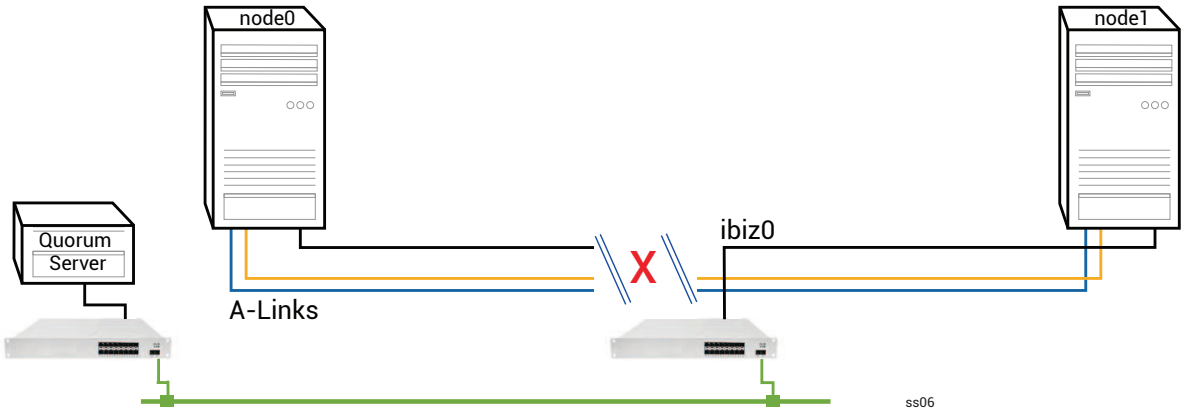
In this SplitSite example, the everRun system includes node0 and node1 with connections identical to those of the system in Example 1. In addition, the system in Example 2 includes a quorum server. The following illustration shows these connections.



A Catastrophic Fault

That careless fork-truck operator crashes through the wall again, severing all of the network connections while leaving the power available and the system running. The following illustration shows the fault

condition.



Fault Handling

The two nodes handle the fault, as follows:

- Node0—The AX on node0 detects the loss of both A-Links as well as all other network paths. Since the node0 AX can no longer detect the presence of its partner, the node0 AX attempts to contact the quorum server. In this case, the quorum server is also unavailable. Therefore, the node0 AX decides to shut down. The shutdown is not a graceful Windows shutdown, but is, instead, an abrupt stop, which causes the application inside the guest VM to stop.
- Node1—The AX on node1 also detects the loss of both A-Links, but ibiz0 remains available. The node1 AX tries to contact the quorum server, which responds, so the node1 AX remains active. The application inside the guest VM runs, perhaps not noticing any problems with the system.

> **Note**: If the node1 AX was not previously active and the guest VM is an HA VM, the guest VM on node1 might need to boot from node1's hard drive. In this case, the application experiences a brief period of downtime while the guest VM boots. (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

After some time, network connectivity is restored: the wall is repaired and the network cables are replaced.

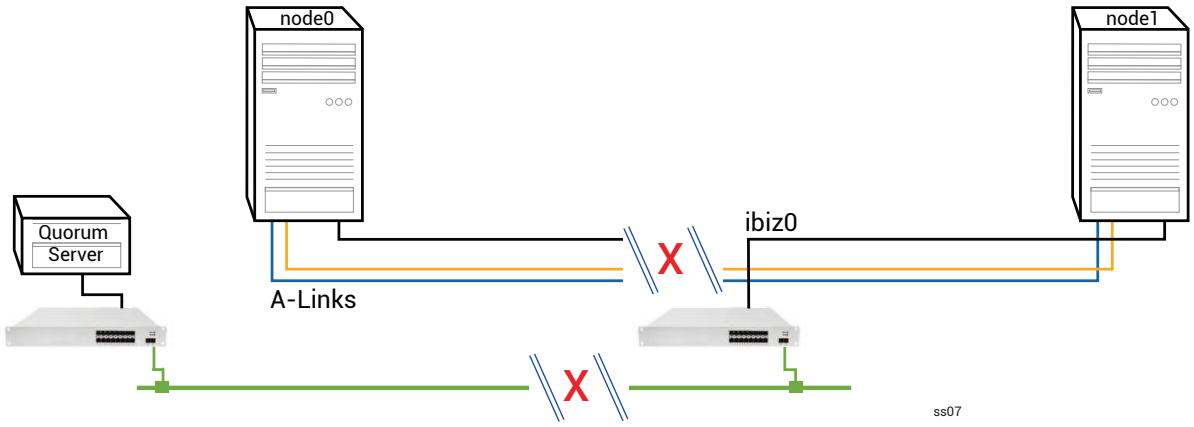
When the node1 AX realizes that its partner is back online, the node0 AX becomes Standby. Because node0 was not previously running, data synchronization begins from node1 to node0.

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 offline.



The fault handling is similar to Example 2 fault handling, with one important difference for node1:

> The node1 AX also detects the loss of both A-Links, but ibiz0 remains available. The node1 AX tries to contact the quorum server, but the communication fails. The AX terminates the guest VM.

In this case, the guest VM is shut down on both node0 and node1, preventing split-brain from occurring. The tradeoff is that the guest VM is unavailable until the connection to either node0 or to the quorum server is restored.

In this case, determine the node that you do not wish to operate and power it down. Then, forcibly boot the node that you wish to operate, and then forcibly boot the VM. For information on shutting down a VM and then starting it, see Managing the Operation of a Virtual Machine.)

Example 2, Modified: The Quorum Server Is Unreachable With No Catastrophic Fault

In some situations, the quorum server might be unreachable even without a catastrophic physical failure. One example is when the quorum computer is rebooted for routine maintenance such as applying an OS patch. In these situations, the AX detects that the quorum service is not responding and so the AX suspends synchronization traffic until the connection to the quorum server is restored. The guest VM continues to run on the node that was active when the connection was lost. However, the guest VM does not move to the standby node because additional faults may occur. After the quorum service is restored, the AX resumes synchronization and normal fault handling, as long as the connection to the quorum server is maintained.

Recovering From a Power Failure

If you restart the system after a power loss or a system shutdown, the 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 **Penguin Solutions 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*nnnnnnn*). You can access the **Penguin Solutions Customer Service Portal** and its Knowledge Base by using your existing portal credentials, or by creating a new user account, as follows.

To access the Knowledge Base

1. Log on to the **Penguin Solutions Customer Service Portal** at https://support.penguinsolutions.com.

   If needed, create a new account as follows:

a. Click **Register**.

b. Enter your contact information including your company email address and registration code, and then click **Submit**.

Your company email address must include a domain name (for example, stratus.com) for a company that is a registered customer of Stratus. The portal sends an email to administrators of the company's account to approve the request.

c. Upon approval, click the link in the email that you receive from Stratus.

d. Enter a new password and finish configuring your account.

If you need assistance creating an account, contact your authorized Penguin Solutions service representative.

2. In the portal, do one of the following:

- In the **Search** box, enter keywords or the KB article number (KB*nnnnnn*) associated with the information you need, and then click the search button.
- Click **Knowledge**, click the name of a product, and then browse available articles.

Related Topics

Supporting Documents

Fixed CVEs

This topic lists Common Vulnerabilities and Exposures (CVE) fixed in the release(s) listed.

CVEs Fixed in everRun Release 8.1.1.0

The following table lists the fixed CVEs (click drop-down icon, if present)

CVEs Fixed in Release 8.1.1.0		
CVE-2020-12351	CVE-2020-12352	CVE-2020-24490
CVE-2021-42780	CVE-2021-42782	CVE-2022-0995
CVE-2022-21546	CVE-2022-28805	CVE-2022-33099
CVE-2022-48703	CVE-2022-48893	CVE-2022-49034

CVEs Fixed in Release 8.1.1.0		
CVE-2022-49063	CVE-2022-49168	CVE-2022-49535
CVE-2022-49636	CVE-2022-49728	CVE-2023-2977
CVE-2023-5992	CVE-2023-28755	CVE-2023-40660
CVE-2023-40661	CVE-2023-52532	CVE-2023-52572
CVE-2023-52621	CVE-2023-52639	CVE-2023-52664
CVE-2023-52757	CVE-2023-52904	CVE-2023-52913
CVE-2023-52917	CVE-2023-52927	CVE-2023-53034
CVE-2023-53188	CVE-2024-3446	CVE-2024-3447
CVE-2024-3567	CVE-2024-4467	CVE-2024-6174
CVE-2024-6505	CVE-2024-7409	CVE-2024-8354
CVE-2024-8443	CVE-2024-8805	CVE-2024-11584
CVE-2024-26686	CVE-2024-26718	CVE-2024-26726
CVE-2024-26739	CVE-2024-26775	CVE-2024-26822
CVE-2024-26837	CVE-2024-26928	CVE-2024-26947
CVE-2024-26982	CVE-2024-27072	CVE-2024-27402
CVE-2024-27407	CVE-2024-35790	CVE-2024-35864
CVE-2024-35866	CVE-2024-35867	CVE-2024-35887
CVE-2024-35904	CVE-2024-35943	CVE-2024-35951
CVE-2024-35963	CVE-2024-35965	CVE-2024-35966

CVEs Fixed in Release 8.1.1.0		
CVE-2024-35967	CVE-2024-36476	CVE-2024-36893
CVE-2024-36899	CVE-2024-36908	CVE-2024-36945
CVE-2024-36968	CVE-2024-38538	CVE-2024-38540
CVE-2024-38541	CVE-2024-38544	CVE-2024-38545
CVE-2024-38553	CVE-2024-38632	CVE-2024-38667
CVE-2024-39463	CVE-2024-39497	CVE-2024-40910
CVE-2024-40953	CVE-2024-40965	CVE-2024-40973
CVE-2024-41016	CVE-2024-41066	CVE-2024-41080
CVE-2024-42079	CVE-2024-42156	CVE-2024-42158
CVE-2024-42230	CVE-2024-42252	CVE-2024-42291
CVE-2024-42315	CVE-2024-43098	CVE-2024-43900
CVE-2024-43904	CVE-2024-44931	CVE-2024-44938
CVE-2024-44939	CVE-2024-44940	CVE-2024-44942
CVE-2024-45615	CVE-2024-45616	CVE-2024-45617
CVE-2024-45618	CVE-2024-45619	CVE-2024-45620
CVE-2024-45828	CVE-2024-46695	CVE-2024-46742
CVE-2024-46751	CVE-2024-46784	CVE-2024-46787
CVE-2024-46809	CVE-2024-46812	CVE-2024-46816
CVE-2024-46821	CVE-2024-46826	CVE-2024-46841

CVEs Fixed in Release 8.1.1.0		
CVE-2024-46849	CVE-2024-46852	CVE-2024-46854
CVE-2024-46855	CVE-2024-46858	CVE-2024-46859
CVE-2024-46865	CVE-2024-46871	CVE-2024-47143
CVE-2024-47408	CVE-2024-47670	CVE-2024-47671
CVE-2024-47672	CVE-2024-47673	CVE-2024-47674
CVE-2024-47679	CVE-2024-47684	CVE-2024-47685
CVE-2024-47690	CVE-2024-47692	CVE-2024-47693
CVE-2024-47695	CVE-2024-47696	CVE-2024-47697
CVE-2024-47698	CVE-2024-47699	CVE-2024-47701
CVE-2024-47705	CVE-2024-47706	CVE-2024-47707
CVE-2024-47709	CVE-2024-47710	CVE-2024-47712
CVE-2024-47713	CVE-2024-47718	CVE-2024-47720
CVE-2024-47723	CVE-2024-47726	CVE-2024-47730
CVE-2024-47734	CVE-2024-47735	CVE-2024-47737
CVE-2024-47739	CVE-2024-47740	CVE-2024-47742
CVE-2024-47747	CVE-2024-47748	CVE-2024-47749
CVE-2024-47756	CVE-2024-47757	CVE-2024-48881
CVE-2024-49571	CVE-2024-49851	CVE-2024-49852
CVE-2024-49856	CVE-2024-49858	CVE-2024-49860

CVEs Fixed in Release 8.1.1.0		
CVE-2024-49863	CVE-2024-49866	CVE-2024-49867
CVE-2024-49868	CVE-2024-49871	CVE-2024-49875
CVE-2024-49877	CVE-2024-49878	CVE-2024-49879
CVE-2024-49881	CVE-2024-49882	CVE-2024-49883
CVE-2024-49884	CVE-2024-49886	CVE-2024-49889
CVE-2024-49890	CVE-2024-49892	CVE-2024-49894
CVE-2024-49895	CVE-2024-49896	CVE-2024-49900
CVE-2024-49902	CVE-2024-49903	CVE-2024-49907
CVE-2024-49913	CVE-2024-49924	CVE-2024-49925
CVE-2024-49927	CVE-2024-49930	CVE-2024-49933
CVE-2024-49935	CVE-2024-49936	CVE-2024-49938
CVE-2024-49944	CVE-2024-49946	CVE-2024-49948
CVE-2024-49949	CVE-2024-49950	CVE-2024-49952
CVE-2024-49954	CVE-2024-49955	CVE-2024-49957
CVE-2024-49958	CVE-2024-49959	CVE-2024-49960
CVE-2024-49962	CVE-2024-49963	CVE-2024-49965
CVE-2024-49966	CVE-2024-49967	CVE-2024-49969
CVE-2024-49973	CVE-2024-49974	CVE-2024-49975
CVE-2024-49977	CVE-2024-49981	CVE-2024-49982

CVEs Fixed in Release 8.1.1.0		
CVE-2024-49983	CVE-2024-49985	CVE-2024-49989
CVE-2024-49995	CVE-2024-49996	CVE-2024-49997
CVE-2024-49998	CVE-2024-50000	CVE-2024-50001
CVE-2024-50002	CVE-2024-50003	CVE-2024-50006
CVE-2024-50007	CVE-2024-50008	CVE-2024-50010
CVE-2024-50013	CVE-2024-50015	CVE-2024-50019
CVE-2024-50024	CVE-2024-50031	CVE-2024-50033
CVE-2024-50035	CVE-2024-50036	CVE-2024-50038
CVE-2024-50039	CVE-2024-50040	CVE-2024-50041
CVE-2024-50044	CVE-2024-50045	CVE-2024-50046
CVE-2024-50047	CVE-2024-50049	CVE-2024-50051
CVE-2024-50055	CVE-2024-50058	CVE-2024-50059
CVE-2024-50062	CVE-2024-50072	CVE-2024-50073
CVE-2024-50074	CVE-2024-50082	CVE-2024-50083
CVE-2024-50085	CVE-2024-50086	CVE-2024-50093
CVE-2024-50095	CVE-2024-50096	CVE-2024-50099
CVE-2024-50101	CVE-2024-50103	CVE-2024-50110
CVE-2024-50115	CVE-2024-50116	CVE-2024-50117
CVE-2024-50121	CVE-2024-50125	CVE-2024-50127

CVEs Fixed in Release 8.1.1.0		
CVE-2024-50128	CVE-2024-50131	CVE-2024-50134
CVE-2024-50141	CVE-2024-50142	CVE-2024-50143
CVE-2024-50148	CVE-2024-50150	CVE-2024-50151
CVE-2024-50153	CVE-2024-50154	CVE-2024-50156
CVE-2024-50160	CVE-2024-50162	CVE-2024-50163
CVE-2024-50167	CVE-2024-50168	CVE-2024-50171
CVE-2024-50179	CVE-2024-50180	CVE-2024-50181
CVE-2024-50182	CVE-2024-50184	CVE-2024-50185
CVE-2024-50186	CVE-2024-50188	CVE-2024-50189
CVE-2024-50191	CVE-2024-50192	CVE-2024-50193
CVE-2024-50194	CVE-2024-50195	CVE-2024-50196
CVE-2024-50198	CVE-2024-50199	CVE-2024-50201
CVE-2024-50202	CVE-2024-50205	CVE-2024-50208
CVE-2024-50209	CVE-2024-50218	CVE-2024-50229
CVE-2024-50230	CVE-2024-50232	CVE-2024-50233
CVE-2024-50234	CVE-2024-50236	CVE-2024-50237
CVE-2024-50242	CVE-2024-50244	CVE-2024-50245
CVE-2024-50247	CVE-2024-50248	CVE-2024-50249
CVE-2024-50251	CVE-2024-50256	CVE-2024-50257

CVEs Fixed in Release 8.1.1.0		
CVE-2024-50258	CVE-2024-50259	CVE-2024-50262
CVE-2024-50264	CVE-2024-50265	CVE-2024-50267
CVE-2024-50268	CVE-2024-50269	CVE-2024-50273
CVE-2024-50275	CVE-2024-50278	CVE-2024-50279
CVE-2024-50280	CVE-2024-50282	CVE-2024-50283
CVE-2024-50287	CVE-2024-50290	CVE-2024-50292
CVE-2024-50295	CVE-2024-50296	CVE-2024-50299
CVE-2024-50301	CVE-2024-50302	CVE-2024-50304
CVE-2024-52332	CVE-2024-52530	CVE-2024-52531
CVE-2024-52532	CVE-2024-53042	CVE-2024-53051
CVE-2024-53052	CVE-2024-53055	CVE-2024-53057
CVE-2024-53058	CVE-2024-53059	CVE-2024-53061
CVE-2024-53063	CVE-2024-53066	CVE-2024-53088
CVE-2024-53096	CVE-2024-53097	CVE-2024-53099
CVE-2024-53101	CVE-2024-53103	CVE-2024-53104
CVE-2024-53112	CVE-2024-53113	CVE-2024-53119
CVE-2024-53120	CVE-2024-53121	CVE-2024-53122
CVE-2024-53124	CVE-2024-53125	CVE-2024-53127
CVE-2024-53128	CVE-2024-53129	CVE-2024-53130

CVEs Fixed in Release 8.1.1.0		
CVE-2024-53131	CVE-2024-53135	CVE-2024-53136
CVE-2024-53138	CVE-2024-53140	CVE-2024-53141
CVE-2024-53142	CVE-2024-53144	CVE-2024-53145
CVE-2024-53146	CVE-2024-53148	CVE-2024-53150
CVE-2024-53151	CVE-2024-53155	CVE-2024-53156
CVE-2024-53157	CVE-2024-53158	CVE-2024-53161
CVE-2024-53164	CVE-2024-53165	CVE-2024-53168
CVE-2024-53171	CVE-2024-53172	CVE-2024-53173
CVE-2024-53174	CVE-2024-53180	CVE-2024-53181
CVE-2024-53183	CVE-2024-53184	CVE-2024-53194
CVE-2024-53197	CVE-2024-53198	CVE-2024-53203
CVE-2024-53206	CVE-2024-53214	CVE-2024-53215
CVE-2024-53217	CVE-2024-53226	CVE-2024-53227
CVE-2024-53237	CVE-2024-53239	CVE-2024-53680
CVE-2024-53685	CVE-2024-53690	CVE-2024-54458
CVE-2024-55881	CVE-2024-55916	CVE-2024-56369
CVE-2024-56531	CVE-2024-56532	CVE-2024-56533
CVE-2024-56539	CVE-2024-56548	CVE-2024-56551
CVE-2024-56558	CVE-2024-56562	CVE-2024-56567

CVEs Fixed in Release 8.1.1.0		
CVE-2024-56568	CVE-2024-56569	CVE-2024-56570
CVE-2024-56572	CVE-2024-56574	CVE-2024-56575
CVE-2024-56576	CVE-2024-56578	CVE-2024-56581
CVE-2024-56586	CVE-2024-56587	CVE-2024-56589
CVE-2024-56590	CVE-2024-56593	CVE-2024-56594
CVE-2024-56595	CVE-2024-56596	CVE-2024-56597
CVE-2024-56598	CVE-2024-56599	CVE-2024-56600
CVE-2024-56601	CVE-2024-56602	CVE-2024-56603
CVE-2024-56605	CVE-2024-56606	CVE-2024-56608
CVE-2024-56610	CVE-2024-56614	CVE-2024-56615
CVE-2024-56616	CVE-2024-56619	CVE-2024-56622
CVE-2024-56623	CVE-2024-56625	CVE-2024-56626
CVE-2024-56627	CVE-2024-56629	CVE-2024-56630
CVE-2024-56631	CVE-2024-56633	CVE-2024-56634
CVE-2024-56636	CVE-2024-56637	CVE-2024-56640
CVE-2024-56642	CVE-2024-56643	CVE-2024-56644
CVE-2024-56645	CVE-2024-56648	CVE-2024-56650
CVE-2024-56651	CVE-2024-56658	CVE-2024-56659
CVE-2024-56662	CVE-2024-56664	CVE-2024-56670

CVEs Fixed in Release 8.1.1.0		
CVE-2024-56672	CVE-2024-56678	CVE-2024-56679
CVE-2024-56681	CVE-2024-56688	CVE-2024-56690
CVE-2024-56691	CVE-2024-56693	CVE-2024-56694
CVE-2024-56698	CVE-2024-56700	CVE-2024-56701
CVE-2024-56704	CVE-2024-56705	CVE-2024-56708
CVE-2024-56715	CVE-2024-56716	CVE-2024-56720
CVE-2024-56721	CVE-2024-56723	CVE-2024-56724
CVE-2024-56726	CVE-2024-56728	CVE-2024-56739
CVE-2024-56745	CVE-2024-56746	CVE-2024-56747
CVE-2024-56748	CVE-2024-56751	CVE-2024-56754
CVE-2024-56756	CVE-2024-56759	CVE-2024-56763
CVE-2024-56767	CVE-2024-56769	CVE-2024-56770
CVE-2024-56774	CVE-2024-56776	CVE-2024-56777
CVE-2024-56778	CVE-2024-56779	CVE-2024-56780
CVE-2024-56781	CVE-2024-56785	CVE-2024-56787
CVE-2024-57791	CVE-2024-57792	CVE-2024-57798
CVE-2024-57802	CVE-2024-57807	CVE-2024-57834
CVE-2024-57838	CVE-2024-57841	CVE-2024-57849
CVE-2024-57850	CVE-2024-57874	CVE-2024-57882

CVEs Fixed in Release 8.1.1.0		
CVE-2024-57883	CVE-2024-57884	CVE-2024-57889
CVE-2024-57890	CVE-2024-57892	CVE-2024-57896
CVE-2024-57897	CVE-2024-57900	CVE-2024-57901
CVE-2024-57902	CVE-2024-57903	CVE-2024-57904
CVE-2024-57906	CVE-2024-57907	CVE-2024-57908
CVE-2024-57910	CVE-2024-57911	CVE-2024-57912
CVE-2024-57913	CVE-2024-57917	CVE-2024-57922
CVE-2024-57925	CVE-2024-57929	CVE-2024-57931
CVE-2024-57938	CVE-2024-57939	CVE-2024-57940
CVE-2024-57946	CVE-2024-57948	CVE-2024-57951
CVE-2024-57973	CVE-2024-57977	CVE-2024-57978
CVE-2024-57979	CVE-2024-57980	CVE-2024-57981
CVE-2024-57986	CVE-2024-57996	CVE-2024-58001
CVE-2024-58002	CVE-2024-58005	CVE-2024-58007
CVE-2024-58010	CVE-2024-58014	CVE-2024-58016
CVE-2024-58017	CVE-2024-58020	CVE-2024-58034
CVE-2024-58051	CVE-2024-58052	CVE-2024-58055
CVE-2024-58058	CVE-2024-58063	CVE-2024-58069
CVE-2024-58071	CVE-2024-58072	CVE-2024-58076

CVEs Fixed in Release 8.1.1.0		
CVE-2024-58079	CVE-2024-58083	CVE-2024-58085
CVE-2024-58086	CVE-2024-58087	CVE-2024-58090
CVE-2024-58093	CVE-2025-0927	CVE-2025-1147
CVE-2025-1148	CVE-2025-1182	CVE-2025-2312
CVE-2025-2784	CVE-2025-3198	CVE-2025-4476
CVE-2025-4945	CVE-2025-4948	CVE-2025-4969
CVE-2025-5244	CVE-2025-5245	CVE-2025-6297
CVE-2025-7545	CVE-2025-7546	CVE-2025-8225
CVE-2025-8961	CVE-2025-9165	CVE-2025-9640
CVE-2025-9900	CVE-2025-10230	CVE-2025-11082
CVE-2025-11083	CVE-2025-20053	CVE-2025-20109
CVE-2025-21090	CVE-2025-21587	CVE-2025-21631
CVE-2025-21636	CVE-2025-21637	CVE-2025-21638
CVE-2025-21639	CVE-2025-21640	CVE-2025-21646
CVE-2025-21647	CVE-2025-21648	CVE-2025-21653
CVE-2025-21664	CVE-2025-21665	CVE-2025-21666
CVE-2025-21669	CVE-2025-21678	CVE-2025-21680
CVE-2025-21683	CVE-2025-21684	CVE-2025-21687
CVE-2025-21689	CVE-2025-21690	CVE-2025-21692

CVEs Fixed in Release 8.1.1.0		
CVE-2025-21694	CVE-2025-21697	CVE-2025-21699
CVE-2025-21700	CVE-2025-21701	CVE-2025-21702
CVE-2025-21703	CVE-2025-21704	CVE-2025-21707
CVE-2025-21708	CVE-2025-21711	CVE-2025-21715
CVE-2025-21718	CVE-2025-21719	CVE-2025-21721
CVE-2025-21722	CVE-2025-21726	CVE-2025-21727
CVE-2025-21728	CVE-2025-21731	CVE-2025-21735
CVE-2025-21736	CVE-2025-21744	CVE-2025-21745
CVE-2025-21748	CVE-2025-21749	CVE-2025-21753
CVE-2025-21756	CVE-2025-21758	CVE-2025-21760
CVE-2025-21761	CVE-2025-21762	CVE-2025-21763
CVE-2025-21764	CVE-2025-21765	CVE-2025-21766
CVE-2025-21767	CVE-2025-21772	CVE-2025-21776
CVE-2025-21779	CVE-2025-21781	CVE-2025-21782
CVE-2025-21785	CVE-2025-21787	CVE-2025-21791
CVE-2025-21795	CVE-2025-21796	CVE-2025-21799
CVE-2025-21802	CVE-2025-21804	CVE-2025-21806
CVE-2025-21811	CVE-2025-21814	CVE-2025-21820
CVE-2025-21823	CVE-2025-21826	CVE-2025-21830

CVEs Fixed in Release 8.1.1.0		
CVE-2025-21835	CVE-2025-21839	CVE-2025-21844
CVE-2025-21846	CVE-2025-21848	CVE-2025-21853
CVE-2025-21858	CVE-2025-21859	CVE-2025-21862
CVE-2025-21865	CVE-2025-21866	CVE-2025-21871
CVE-2025-21875	CVE-2025-21877	CVE-2025-21878
CVE-2025-21887	CVE-2025-21888	CVE-2025-21898
CVE-2025-21904	CVE-2025-21905	CVE-2025-21909
CVE-2025-21910	CVE-2025-21912	CVE-2025-21914
CVE-2025-21916	CVE-2025-21917	CVE-2025-21919
CVE-2025-21920	CVE-2025-21922	CVE-2025-21924
CVE-2025-21925	CVE-2025-21926	CVE-2025-21928
CVE-2025-21934	CVE-2025-21935	CVE-2025-21941
CVE-2025-21943	CVE-2025-21948	CVE-2025-21950
CVE-2025-21951	CVE-2025-21956	CVE-2025-21957
CVE-2025-21959	CVE-2025-21962	CVE-2025-21963
CVE-2025-21964	CVE-2025-21968	CVE-2025-21970
CVE-2025-21971	CVE-2025-21975	CVE-2025-21981
CVE-2025-21991	CVE-2025-21992	CVE-2025-21993
CVE-2025-21996	CVE-2025-21999	CVE-2025-22004

CVEs Fixed in Release 8.1.1.0		
CVE-2025-22005	CVE-2025-22007	CVE-2025-22008
CVE-2025-22010	CVE-2025-22014	CVE-2025-22018
CVE-2025-22020	CVE-2025-22021	CVE-2025-22025
CVE-2025-22027	CVE-2025-22035	CVE-2025-22044
CVE-2025-22045	CVE-2025-22050	CVE-2025-22054
CVE-2025-22055	CVE-2025-22056	CVE-2025-22060
CVE-2025-22062	CVE-2025-22063	CVE-2025-22066
CVE-2025-22071	CVE-2025-22073	CVE-2025-22075
CVE-2025-22079	CVE-2025-22081	CVE-2025-22086
CVE-2025-22089	CVE-2025-22097	CVE-2025-22839
CVE-2025-22840	CVE-2025-22889	CVE-2025-23136
CVE-2025-23138	CVE-2025-23140	CVE-2025-23142
CVE-2025-23144	CVE-2025-23145	CVE-2025-23146
CVE-2025-23147	CVE-2025-23150	CVE-2025-23151
CVE-2025-23156	CVE-2025-23157	CVE-2025-23158
CVE-2025-23159	CVE-2025-23161	CVE-2025-24294
CVE-2025-24305	CVE-2025-26403	CVE-2025-30691
CVE-2025-30698	CVE-2025-30749	CVE-2025-30754
CVE-2025-32050	CVE-2025-32051	CVE-2025-32052

CVEs Fixed in Release 8.1.1.0		
CVE-2025-32053	CVE-2025-32086	CVE-2025-32906
CVE-2025-32907	CVE-2025-32908	CVE-2025-32909
CVE-2025-32910	CVE-2025-32911	CVE-2025-32912
CVE-2025-32913	CVE-2025-32914	CVE-2025-37738
CVE-2025-37739	CVE-2025-37740	CVE-2025-37742
CVE-2025-37749	CVE-2025-37752	CVE-2025-37756
CVE-2025-37758	CVE-2025-37765	CVE-2025-37766
CVE-2025-37767	CVE-2025-37768	CVE-2025-37770
CVE-2025-37771	CVE-2025-37773	CVE-2025-37780
CVE-2025-37781	CVE-2025-37785	CVE-2025-37787
CVE-2025-37788	CVE-2025-37789	CVE-2025-37790
CVE-2025-37792	CVE-2025-37794	CVE-2025-37796
CVE-2025-37797	CVE-2025-37798	CVE-2025-37803
CVE-2025-37805	CVE-2025-37808	CVE-2025-37810
CVE-2025-37811	CVE-2025-37812	CVE-2025-37817
CVE-2025-37819	CVE-2025-37823	CVE-2025-37824
CVE-2025-37829	CVE-2025-37830	CVE-2025-37836
CVE-2025-37838	CVE-2025-37839	CVE-2025-37840
CVE-2025-37841	CVE-2025-37844	CVE-2025-37850

CVEs Fixed in Release 8.1.1.0		
CVE-2025-37857	CVE-2025-37858	CVE-2025-37859
CVE-2025-37862	CVE-2025-37867	CVE-2025-37875
CVE-2025-37881	CVE-2025-37883	CVE-2025-37889
CVE-2025-37890	CVE-2025-37892	CVE-2025-37905
CVE-2025-37909	CVE-2025-37911	CVE-2025-37912
CVE-2025-37913	CVE-2025-37914	CVE-2025-37923
CVE-2025-37927	CVE-2025-37930	CVE-2025-37932
CVE-2025-37937	CVE-2025-37940	CVE-2025-37948
CVE-2025-37949	CVE-2025-37958	CVE-2025-37963
CVE-2025-37964	CVE-2025-37967	CVE-2025-37969
CVE-2025-37970	CVE-2025-37982	CVE-2025-37983
CVE-2025-37985	CVE-2025-37989	CVE-2025-37991
CVE-2025-37992	CVE-2025-37994	CVE-2025-37995
CVE-2025-37997	CVE-2025-37998	CVE-2025-38000
CVE-2025-38001	CVE-2025-38003	CVE-2025-38004
CVE-2025-38005	CVE-2025-38009	CVE-2025-38023
CVE-2025-38024	CVE-2025-38031	CVE-2025-38034
CVE-2025-38035	CVE-2025-38037	CVE-2025-38043
CVE-2025-38044	CVE-2025-38048	CVE-2025-38051

CVEs Fixed in Release 8.1.1.0		
CVE-2025-38052	CVE-2025-38058	CVE-2025-38061
CVE-2025-38065	CVE-2025-38066	CVE-2025-38067
CVE-2025-38068	CVE-2025-38072	CVE-2025-38074
CVE-2025-38075	CVE-2025-38077	CVE-2025-38078
CVE-2025-38079	CVE-2025-38083	CVE-2025-38084
CVE-2025-38085	CVE-2025-38086	CVE-2025-38088
CVE-2025-38090	CVE-2025-38100	CVE-2025-38102
CVE-2025-38103	CVE-2025-38107	CVE-2025-38108
CVE-2025-38111	CVE-2025-38112	CVE-2025-38115
CVE-2025-38119	CVE-2025-38120	CVE-2025-38122
CVE-2025-38135	CVE-2025-38136	CVE-2025-38138
CVE-2025-38143	CVE-2025-38145	CVE-2025-38146
CVE-2025-38147	CVE-2025-38152	CVE-2025-38153
CVE-2025-38154	CVE-2025-38157	CVE-2025-38159
CVE-2025-38160	CVE-2025-38161	CVE-2025-38163
CVE-2025-38167	CVE-2025-38173	CVE-2025-38174
CVE-2025-38181	CVE-2025-38184	CVE-2025-38194
CVE-2025-38197	CVE-2025-38200	CVE-2025-38203
CVE-2025-38204	CVE-2025-38206	CVE-2025-38211

CVEs Fixed in Release 8.1.1.0		
CVE-2025-38212	CVE-2025-38218	CVE-2025-38219
CVE-2025-38222	CVE-2025-38226	CVE-2025-38227
CVE-2025-38229	CVE-2025-38231	CVE-2025-38237
CVE-2025-38245	CVE-2025-38249	CVE-2025-38251
CVE-2025-38257	CVE-2025-38262	CVE-2025-38263
CVE-2025-38273	CVE-2025-38280	CVE-2025-38285
CVE-2025-38286	CVE-2025-38293	CVE-2025-38298
CVE-2025-38305	CVE-2025-38310	CVE-2025-38312
CVE-2025-38313	CVE-2025-38319	CVE-2025-38320
CVE-2025-38324	CVE-2025-38326	CVE-2025-38328
CVE-2025-38332	CVE-2025-38336	CVE-2025-38337
CVE-2025-38342	CVE-2025-38344	CVE-2025-38345
CVE-2025-38346	CVE-2025-38348	CVE-2025-38350
CVE-2025-38352	CVE-2025-38362	CVE-2025-38363
CVE-2025-38371	CVE-2025-38375	CVE-2025-38377
CVE-2025-38384	CVE-2025-38386	CVE-2025-38387
CVE-2025-38389	CVE-2025-38391	CVE-2025-38393
CVE-2025-38395	CVE-2025-38399	CVE-2025-38400
CVE-2025-38401	CVE-2025-38403	CVE-2025-38406

CVEs Fixed in Release 8.1.1.0		
CVE-2025-38410	CVE-2025-38412	CVE-2025-38415
CVE-2025-38416	CVE-2025-38418	CVE-2025-38419
CVE-2025-38420	CVE-2025-38424	CVE-2025-38428
CVE-2025-38430	CVE-2025-38439	CVE-2025-38441
CVE-2025-38443	CVE-2025-38444	CVE-2025-38445
CVE-2025-38448	CVE-2025-38457	CVE-2025-38458
CVE-2025-38459	CVE-2025-38460	CVE-2025-38461
CVE-2025-38462	CVE-2025-38464	CVE-2025-38465
CVE-2025-38466	CVE-2025-38467	CVE-2025-38477
CVE-2025-38498	CVE-2025-38513	CVE-2025-38514
CVE-2025-38515	CVE-2025-38516	CVE-2025-38540
CVE-2025-38542	CVE-2025-38575	CVE-2025-38618
CVE-2025-39728	CVE-2025-39735	CVE-2025-39889
CVE-2025-41244	CVE-2025-46420	CVE-2025-46421
CVE-2025-48989	CVE-2025-50059	CVE-2025-50106
CVE-2025-52434	CVE-2025-52520	CVE-2025-53057
CVE-2025-53066	CVE-2025-53506	CVE-2025-55752
CVE-2025-58060	CVE-2025-58364	

CVEs Fixed in everRun Release 8.1.0.2

No CVEs fixed.

CVEs Fixed in everRun Release 8.1.0.1

The following table lists the fixed CVEs (click drop-down icon, if present)

CVEs Fixed in Release 8.1.0.1		
CVE-2022-49043	CVE-2023-25193	CVE-2023-48795
CVE-2024-3596	CVE-2024-8176	CVE-2024-9143
CVE-2024-11053	CVE-2024-11168	CVE-2024-11187
CVE-2024-12705	CVE-2024-13176	CVE-2024-21820
CVE-2024-21853	CVE-2024-23918	CVE-2024-23984
CVE-2024-24968	CVE-2024-34459	CVE-2024-35176
CVE-2024-39908	CVE-2024-41123	CVE-2024-43398
CVE-2024-48916	CVE-2024-50349	CVE-2024-50602
CVE-2024-52006	CVE-2024-55549	CVE-2024-56171
CVE-2024-56201	CVE-2024-56326	CVE-2024-56406
CVE-2025-0395	CVE-2025-0938	CVE-2025-1795
CVE-2025-3277	CVE-2025-4516	CVE-2025-4802
CVE-2025-5054	CVE-2025-5914	CVE-2025-5915
CVE-2025-5916	CVE-2025-5917	CVE-2025-6019
CVE-2025-6020	CVE-2025-6199	CVE-2025-6965
CVE-2025-7345	CVE-2025-24855	CVE-2025-24928
CVE-2025-25186	CVE-2025-27113	CVE-2025-27219

CVEs Fixed in Release 8.1.0.1		
CVE-2025-27220	CVE-2025-27221	CVE-2025-27363
CVE-2025-27516	CVE-2025-27613	CVE-2025-27614
CVE-2025-29087	CVE-2025-29088	CVE-2025-30258
CVE-2025-32414	CVE-2025-32415	CVE-2025-32462
CVE-2025-32463	CVE-2025-46835	CVE-2025-47273
CVE-2025-48384	CVE-2025-48385	CVE-2025-48386

CVEs Fixed in everRun Release 8.1.0.0

The following table lists the fixed CVEs (click drop-down icon, if present)

CVEs Fixed in Release 8.1.0.0		
CVE-2016-1585	CVE-2016-9840	CVE-2016-9841
CVE-2017-16516	CVE-2018-25032	CVE-2018-25091
CVE-2019-9511	CVE-2019-9513	CVE-2020-14394
CVE-2020-24165	CVE-2021-3611	CVE-2021-3638
CVE-2021-32292	CVE-2021-40633	CVE-2022-3715
CVE-2022-24795	CVE-2022-28506	CVE-2022-37434
CVE-2022-45939	CVE-2022-46908	CVE-2022-47695
CVE-2022-48063	CVE-2022-48065	CVE-2022-48174
CVE-2022-48337	CVE-2022-48338	CVE-2022-48339
CVE-2022-48340	CVE-2022-48522	CVE-2022-48622

CVEs Fixed in Release 8.1.0.0		
CVE-2023-2953	CVE-2023-3180	CVE-2023-3255
CVE-2023-3301	CVE-2023-3354	CVE-2023-4135
CVE-2023-4156	CVE-2023-5088	CVE-2023-5981
CVE-2023-7104	CVE-2023-28617	CVE-2023-31315
CVE-2023-33460	CVE-2023-34058	CVE-2023-34059
CVE-2023-39742	CVE-2023-39804	CVE-2023-40360
CVE-2023-42363	CVE-2023-42364	CVE-2023-42365
CVE-2023-42467	CVE-2023-42667	CVE-2023-43040
CVE-2023-43785	CVE-2023-43786	CVE-2023-43787
CVE-2023-43804	CVE-2023-44487	CVE-2023-45803
CVE-2023-47038	CVE-2023-48795	CVE-2023-49141
CVE-2023-50782	CVE-2023-52889	CVE-2023-52918
CVE-2024-0760	CVE-2024-1737	CVE-2024-1975
CVE-2024-3651	CVE-2024-4076	CVE-2024-6119
CVE-2024-6345	CVE-2024-6387	CVE-2024-6655
CVE-2024-7006	CVE-2024-9287	CVE-2024-9681
CVE-2024-10224	CVE-2024-11003	CVE-2024-12084
CVE-2024-12085	CVE-2024-12086	CVE-2024-12087
CVE-2024-12088	CVE-2024-12747	CVE-2024-20696

CVEs Fixed in Release 8.1.0.0		
CVE-2024-21131	CVE-2024-21138	CVE-2024-21140
CVE-2024-21145	CVE-2024-21147	CVE-2024-21208
CVE-2024-21210	CVE-2024-21217	CVE-2024-21235
CVE-2024-22365	CVE-2024-22667	CVE-2024-24806
CVE-2024-24853	CVE-2024-24980	CVE-2024-25062
CVE-2024-25744	CVE-2024-25939	CVE-2024-26130
CVE-2024-26607	CVE-2024-26661	CVE-2024-26669
CVE-2024-26800	CVE-2024-26893	CVE-2024-28085
CVE-2024-28182	CVE-2024-30203	CVE-2024-30204
CVE-2024-30205	CVE-2024-32002	CVE-2024-32004
CVE-2024-32020	CVE-2024-32021	CVE-2024-32465
CVE-2024-32487	CVE-2024-33599	CVE-2024-33600
CVE-2024-33601	CVE-2024-33602	CVE-2024-35176
CVE-2024-36484	CVE-2024-37370	CVE-2024-37371
CVE-2024-38428	CVE-2024-38577	CVE-2024-38602
CVE-2024-38611	CVE-2024-39331	CVE-2024-39472
CVE-2024-39908	CVE-2024-40897	CVE-2024-40915
CVE-2024-41011	CVE-2024-41012	CVE-2024-41015
CVE-2024-41017	CVE-2024-41019	CVE-2024-41020

CVEs Fixed in Release 8.1.0.0		
CVE-2024-41022	CVE-2024-41042	CVE-2024-41059
CVE-2024-41060	CVE-2024-41063	CVE-2024-41064
CVE-2024-41065	CVE-2024-41068	CVE-2024-41070
CVE-2024-41071	CVE-2024-41072	CVE-2024-41073
CVE-2024-41077	CVE-2024-41078	CVE-2024-41081
CVE-2024-41090	CVE-2024-41091	CVE-2024-41098
CVE-2024-41123	CVE-2024-41671	CVE-2024-41946
CVE-2024-42114	CVE-2024-42126	CVE-2024-42246
CVE-2024-42259	CVE-2024-42265	CVE-2024-42267
CVE-2024-42269	CVE-2024-42270	CVE-2024-42271
CVE-2024-42272	CVE-2024-42274	CVE-2024-42276
CVE-2024-42277	CVE-2024-42280	CVE-2024-42281
CVE-2024-42283	CVE-2024-42284	CVE-2024-42285
CVE-2024-42286	CVE-2024-42287	CVE-2024-42288
CVE-2024-42289	CVE-2024-42290	CVE-2024-42292
CVE-2024-42295	CVE-2024-42296	CVE-2024-42297
CVE-2024-42299	CVE-2024-42301	CVE-2024-42302
CVE-2024-42304	CVE-2024-42305	CVE-2024-42306
CVE-2024-42309	CVE-2024-42310	CVE-2024-42311

CVEs Fixed in Release 8.1.0.0		
CVE-2024-42312	CVE-2024-42313	CVE-2024-42318
CVE-2024-43817	CVE-2024-43828	CVE-2024-43829
CVE-2024-43830	CVE-2024-43834	CVE-2024-43835
CVE-2024-43839	CVE-2024-43841	CVE-2024-43846
CVE-2024-43849	CVE-2024-43853	CVE-2024-43854
CVE-2024-43856	CVE-2024-43858	CVE-2024-43860
CVE-2024-43861	CVE-2024-43863	CVE-2024-43867
CVE-2024-43869	CVE-2024-43870	CVE-2024-43871
CVE-2024-43873	CVE-2024-43875	CVE-2024-43879
CVE-2024-43880	CVE-2024-43882	CVE-2024-43883
CVE-2024-43884	CVE-2024-43889	CVE-2024-43890
CVE-2024-43892	CVE-2024-43893	CVE-2024-43894
CVE-2024-43902	CVE-2024-43905	CVE-2024-43907
CVE-2024-43908	CVE-2024-43909	CVE-2024-43914
CVE-2024-44934	CVE-2024-44935	CVE-2024-44944
CVE-2024-44946	CVE-2024-44947	CVE-2024-44948
CVE-2024-44954	CVE-2024-44958	CVE-2024-44960
CVE-2024-44965	CVE-2024-44966	CVE-2024-44969
CVE-2024-44971	CVE-2024-44974	CVE-2024-44982

CVEs Fixed in Release 8.1.0.0		
CVE-2024-44983	CVE-2024-44985	CVE-2024-44986
CVE-2024-44987	CVE-2024-44988	CVE-2024-44989
CVE-2024-44990	CVE-2024-44995	CVE-2024-44998
CVE-2024-44999	CVE-2024-45003	CVE-2024-45006
CVE-2024-45007	CVE-2024-45008	CVE-2024-45009
CVE-2024-45011	CVE-2024-45018	CVE-2024-45021
CVE-2024-45025	CVE-2024-45026	CVE-2024-45028
CVE-2024-45490	CVE-2024-45491	CVE-2024-45492
CVE-2024-46673	CVE-2024-46675	CVE-2024-46676
CVE-2024-46677	CVE-2024-46679	CVE-2024-46685
CVE-2024-46689	CVE-2024-46702	CVE-2024-46707
CVE-2024-46713	CVE-2024-46714	CVE-2024-46719
CVE-2024-46721	CVE-2024-46722	CVE-2024-46723
CVE-2024-46724	CVE-2024-46725	CVE-2024-46731
CVE-2024-46732	CVE-2024-46737	CVE-2024-46738
CVE-2024-46739	CVE-2024-46740	CVE-2024-46743
CVE-2024-46744	CVE-2024-46745	CVE-2024-46746
CVE-2024-46747	CVE-2024-46750	CVE-2024-46752
CVE-2024-46755	CVE-2024-46756	CVE-2024-46757

CVEs Fixed in Release 8.1.0.0		
CVE-2024-46758	CVE-2024-46759	CVE-2024-46761
CVE-2024-46763	CVE-2024-46771	CVE-2024-46777
CVE-2024-46780	CVE-2024-46781	CVE-2024-46782
CVE-2024-46783	CVE-2024-46791	CVE-2024-46795
CVE-2024-46798	CVE-2024-46800	CVE-2024-46804
CVE-2024-46805	CVE-2024-46807	CVE-2024-46810
CVE-2024-46814	CVE-2024-46815	CVE-2024-46817
CVE-2024-46818	CVE-2024-46819	CVE-2024-46822
CVE-2024-46828	CVE-2024-46829	CVE-2024-46832
CVE-2024-46840	CVE-2024-46844	CVE-2024-47175
CVE-2024-47220	CVE-2024-47538	CVE-2024-47541
CVE-2024-47542	CVE-2024-47600	CVE-2024-47606
CVE-2024-47607	CVE-2024-47615	CVE-2024-47659
CVE-2024-47660	CVE-2024-47663	CVE-2024-47665
CVE-2024-47667	CVE-2024-47668	CVE-2024-47669
CVE-2024-47835	CVE-2024-48990	CVE-2024-48991
CVE-2024-48992	CVE-2024-49761	CVE-2024-52530
CVE-2024-52531	CVE-2024-52532	CVE-2024-52533

CVEs Fixed in everRun Release 8.0.0.0

The following table lists the fixed CVEs (click drop-down icon, if present)

CVEs Fixed in Release 8.0.0.0		
CVE-2020-22218	CVE-2020-22218	CVE-2021-40211
CVE-2021-4104	CVE-2022-3564	CVE-2022-40982
CVE-2022-40982	CVE-2022-48339	CVE-2023-1999
CVE-2023-21930	CVE-2023-21937	CVE-2023-21938
CVE-2023-21939	CVE-2023-21954	CVE-2023-21967
CVE-2023-21968	CVE-2023-24329	CVE-2023-24329
CVE-2023-2828	CVE-2023-32067	CVE-2023-32233
CVE-2023-35001	CVE-2023-3609	CVE-2023-3611
CVE-2023-3611	CVE-2023-3776	CVE-2023-3776
CVE-2023-38408	CVE-2023-4206	CVE-2023-4206
CVE-2023-4207	CVE-2023-4207	CVE-2023-4208
CVE-2023-4208	CVE-2023-42753	CVE-2023-44446
CVE-2023-6377	CVE-2023-6478	CVE-2023-6816
CVE-2024-0229	CVE-2024-20918	CVE-2024-20919
CVE-2024-20921	CVE-2024-20926	CVE-2024-20945
CVE-2024-20952	CVE-2024-21885	CVE-2024-21886

CVEs Fixed in everRun Release 7.9.4.1

No CVEs fixed.

CVEs Fixed in everRun Release 7.9.4.0

The following table lists the fixed CVEs (click drop-down icon, if present)

CVEs Fixed in Release 7.9.4.0		
CVE-2020-22218	CVE-2021-40211	CVE-2022-3564
CVE-2022-40982	CVE-2022-42896	CVE-2022-48339
CVE-2023-1999	CVE-2023-2828	CVE-2023-3341
CVE-2023-3609	CVE-2023-3611	CVE-2023-3776
CVE-2023-4206	CVE-2023-4207	CVE-2023-4208
CVE-2023-4921	CVE-2023-6377	CVE-2023-6478
CVE-2023-6816	CVE-2023-21930	CVE-2023-21937
CVE-2023-21938	CVE-2023-21939	CVE-2023-21954
CVE-2023-21967	CVE-2023-21968	CVE-2023-24329
CVE-2023-32067	CVE-2023-32233	CVE-2023-35001
CVE-2023-38408	CVE-2023-38409	CVE-2023-42753
CVE-2023-44446	CVE-2023-45871	CVE-2024-0229
CVE-2024-1086	CVE-2024-20918	CVE-2024-20919
CVE-2024-20921	CVE-2024-20926	CVE-2024-20945
CVE-2024-20952	CVE-2024-21885	CVE-2024-21886
CVE-2024-26602		

CVEs Fixed in everRun Release 7.9.3.1

No CVEs fixed.

CVEs Fixed in everRun Release 7.9.3.0

The following table lists the fixed CVEs (click drop-down icon, if present)

CVEs Fixed in Release 7.9.3.0		
CVE-2013-2566	CVE-2015-2808	CVE-2017-5715
CVE-2018-13405	CVE-2020-0465	CVE-2020-0466
CVE-2020-25717	CVE-2020-36385	CVE-2021-0920
CVE-2021-3564	CVE-2021-3573	CVE-2021-3752
CVE-2021-4037	CVE-2021-4155	CVE-2021-20271
CVE-2021-25220	CVE-2021-26401	CVE-2021-41617
CVE-2021-44142	CVE-2022-0330	CVE-2022-0492
CVE-2022-1729	CVE-2022-1966	CVE-2022-2526
CVE-2022-2795	CVE-2022-2964	CVE-2022-4254
CVE-2022-4283	CVE-2022-4378	CVE-2022-4883
CVE-2022-22942	CVE-2022-24407	CVE-2022-29154
CVE-2022-32250	CVE-2022-37434	CVE-2022-38023
CVE-2022-38177	CVE-2022-38178	CVE-2022-40674
CVE-2022-41974	CVE 2022-42703	CVE-2022-42898
CVE-2022-46340	CVE-2022-46341	CVE-2022-46342
CVE-2022-46343	CVE-2022-46344	CVE-2023-0286
CVE-2023-0494	CVE-2023-0767	CVE-2023-22809

CVEs Fixed in everRun Release 7.9.2.0

The following table lists the fixed CVEs (click drop-down icon, if present)

CVEs Fixed in Release 7.9.2.0		
CVE-2018-25032	CVE-2020-25709	CVE-2020-25710
CVE-2020-26116	CVE-2020-26137	CVE-2021-3177
CVE-2021-45960	CVE-2021-46143	CVE-2022-0778
CVE-2022-1271	CVE-2022-21426	CVE-2022-21434
CVE-2022-21443	CVE-2022-21476	CVE-2022-21496
CVE-2022-21540	CVE-2022-21541	CVE-2022-22822
CVE-2022-22823	CVE-2022-22824	CVE-2022-22825
CVE-2022-22826	CVE-2022-22827	CVE-2022-23852
CVE-2022-25235	CVE-2022-25236	CVE-2022-25315
CVE-2022-34169		

CVEs Fixed in everRun Release 7.9.1.1

No CVEs fixed.

CVEs Fixed in everRun Release 7.9.1.0

The following table lists the fixed CVEs (click drop-down icon, if present)

CVEs Fixed in Release 7.9.1.0		
CVE-2016-2124	CVE-2016-4658	CVE-2018-25011
CVE-2019-20934	CVE-2020-0543	CVE-2020-0548
CVE-2020-0549	CVE-2020-8648	CVE-2020-8695
CVE-2020-8696	CVE-2020-8698	CVE-2020-11668

CVEs Fixed in Release 7.9.1.0		
CVE-2020-12362	CVE-2020-12363	CVE-2020-12364
CVE-2020-24489	CVE-2020-24511	CVE-2020-24512
CVE-2020-24513	CVE-2020-25717	CVE-2020-27170
CVE-2020-27777	CVE-2020-29443	CVE-2020-36328
CVE-2020-36329	CVE-2021-2341	CVE-2021-2369
CVE-2021-2388	CVE-2021-3246	CVE-2021-3347
CVE-2021-3472	CVE-2021-3621	CVE-2021-3653
CVE-2021-3656	CVE-2021-3715	CVE-2021-4034
CVE-2021-20254	CVE-2021-22543	CVE-2021-22555
CVE-2021-23840	CVE-2021-23841	CVE-2021-25214
CVE-2021-27219	CVE-2021-29154	CVE-2021-29650
CVE-2021-31535	CVE-2021-32027	CVE-2021-32399
CVE-2021-33033	CVE-2021-33034	CVE-2021-33909
CVE-2021-35550	CVE-2021-35556	CVE-2021-35559
CVE-2021-35561	CVE-2021-35564	CVE-2021-35565
CVE-2021-35567	CVE-2021-35578	CVE-2021-35586
CVE-2021-35588	CVE-2021-35603	CVE-2021-37576
CVE-2021-42574	CVE-2021-43527	

CVEs Fixed in everRun Release 7.9.0.0

The following table lists the fixed CVEs (click drop-down icon, if present)

CVEs Fixed in Release 7.9.0.0		
CVE-2013-2139	CVE-2015-2716	CVE-2015-6360
CVE-2016-5766	CVE-2017-12652	CVE-2017-15715
CVE-2017-18190	CVE-2017-18551	CVE-2018-1283
CVE-2018-1303	CVE-2018-11782	CVE-2018-15746
CVE-2018-19662	CVE-2018-20836	CVE-2018-20843
CVE-2019-2974	CVE-2019-5094	CVE-2019-5188
CVE-2019-5482	CVE-2019-6237	CVE-2019-6251
CVE-2019-6978	CVE-2019-7572	CVE-2019-7573
CVE-2019-7574	CVE-2019-7575	CVE-2019-7576
CVE-2019-7577	CVE-2019-7578	CVE-2019-7635
CVE-2019-7636	CVE-2019-7637	CVE-2019-7638
CVE-2019-8506	CVE-2019-8524	CVE-2019-8535
CVE-2019-8536	CVE-2019-8544	CVE-2019-8551
CVE-2019-8558	CVE-2019-8559	CVE-2019-8563
CVE-2019-8571	CVE-2019-8583	CVE-2019-8584
CVE-2019-8586	CVE-2019-8587	CVE-2019-8594
CVE-2019-8595	CVE-2019-8596	CVE-2019-8597
CVE-2019-8601	CVE-2019-8607	CVE-2019-8608
CVE-2019-8609	CVE-2019-8610	CVE-2019-8611

CVEs Fixed in Release 7.9.0.0		
CVE-2019-8615	CVE-2019-8619	CVE-2019-8622
CVE-2019-8623	CVE-2019-8625	CVE-2019-8644
CVE-2019-8649	CVE-2019-8658	CVE-2019-8666
CVE-2019-8669	CVE-2019-8671	CVE-2019-8672
CVE-2019-8673	CVE-2019-8674	CVE-2019-8675
CVE-2019-8676	CVE-2019-8677	CVE-2019-8678
CVE-2019-8679	CVE-2019-8680	CVE-2019-8681
CVE-2019-8683	CVE-2019-8684	CVE-2019-8686
CVE-2019-8687	CVE-2019-8688	CVE-2019-8689
CVE-2019-8690	CVE-2019-8696	CVE-2019-8707
CVE-2019-8710	CVE-2019-8719	CVE-2019-8720
CVE-2019-8726	CVE-2019-8733	CVE-2019-8735
CVE-2019-8743	CVE-2019-8763	CVE-2019-8764
CVE-2019-8765	CVE-2019-8766	CVE-2019-8768
CVE-2019-8769	CVE-2019-8771	CVE-2019-8782
CVE-2019-8783	CVE-2019-8808	CVE-2019-8811
CVE-2019-8812	CVE-2019-8813	CVE-2019-8814
CVE-2019-8815	CVE-2019-8816	CVE-2019-8819
CVE-2019-8820	CVE-2019-8821	CVE-2019-8822

CVEs Fixed in Release 7.9.0.0		
CVE-2019-8823	CVE-2019-8835	CVE-2019-8844
CVE-2019-8846	CVE-2019-9454	CVE-2019-9458
CVE-2019-10098	CVE-2019-10208	CVE-2019-11068
CVE-2019-11070	CVE-2019-11719	CVE-2019-11727
CVE-2019-11756	CVE-2019-12450	CVE-2019-12614
CVE-2019-12749	CVE-2019-14494	CVE-2019-14744
CVE-2019-14822	CVE-2019-14834	CVE-2019-14866
CVE-2019-14907	CVE-2019-14973	CVE-2019-15217
CVE-2019-15691	CVE-2019-15692	CVE-2019-15693
CVE-2019-15694	CVE-2019-15695	CVE-2019-15807
CVE-2019-15903	CVE-2019-15917	CVE-2019-16231
CVE-2019-16233	CVE-2019-16707	CVE-2019-16935
CVE-2019-16994	CVE-2019-17006	CVE-2019-17023
CVE-2019-17053	CVE-2019-17055	CVE-2019-17498
CVE-2019-17546	CVE-2019-17563	CVE-2019-18197
CVE-2019-18282	CVE-2019-18808	CVE-2019-19046
CVE-2019-19055	CVE-2019-19058	CVE-2019-19059
CVE-2019-19062	CVE-2019-19063	CVE-2019-19126
CVE-2019-19332	CVE-2019-19447	CVE-2019-19523

CVEs Fixed in Release 7.9.0.0		
CVE-2019-19524	CVE-2019-19530	CVE-2019-19532
CVE-2019-19534	CVE-2019-19537	CVE-2019-19767
CVE-2019-19807	CVE-2019-19956	CVE-2019-20054
CVE-2019-20095	CVE-2019-20382	CVE-2019-20386
CVE-2019-20388	CVE-2019-20485	CVE-2019-20636
CVE-2019-20811	CVE-2019-20907	CVE-2019-25013
CVE-2020-0427	CVE-2020-1472	CVE-2020-1749
CVE-2020-1927	CVE-2020-1934	CVE-2020-1935
CVE-2020-1971	CVE-2020-1983	CVE-2020-2574
CVE-2020-2732	CVE-2020-2752	CVE-2020-2780
CVE-2020-2812	CVE-2020-3862	CVE-2020-3864
CVE-2020-3865	CVE-2020-3867	CVE-2020-3868
CVE-2020-3885	CVE-2020-3894	CVE-2020-3895
CVE-2020-3897	CVE-2020-3899	CVE-2020-3900
CVE-2020-3901	CVE-2020-3902	CVE-2020-5313
CVE-2020-6829	CVE-2020-7053	CVE-2020-7595
CVE-2020-8177	CVE-2020-8622	CVE-2020-8623
CVE-2020-8624	CVE-2020-8625	CVE-2020-8647
CVE-2020-8649	CVE-2020-8695	CVE-2020-8696

CVEs Fixed in Release 7.9.0.0		
CVE-2020-8698	CVE-2020-9383	CVE-2020-10018
CVE-2020-10029	CVE-2020-10543	CVE-2020-10690
CVE-2020-10703	CVE-2020-10713	CVE-2020-10732
CVE-2020-10742	CVE-2020-10751	CVE-2020-10754
CVE-2020-10769	CVE-2020-10878	CVE-2020-10942
CVE-2020-11078	CVE-2020-11565	CVE-2020-11761
CVE-2020-11763	CVE-2020-11764	CVE-2020-11793
CVE-2020-12243	CVE-2020-12321	CVE-2020-12351
CVE-2020-12352	CVE-2020-12400	CVE-2020-12401
CVE-2020-12402	CVE-2020-12403	CVE-2020-12723
CVE-2020-12770	CVE-2020-12825	CVE-2020-12826
CVE-2020-13765	CVE-2020-13935	CVE-2020-14305
CVE-2020-14308	CVE-2020-14309	CVE-2020-14310
CVE-2020-14311	CVE-2020-14314	CVE-2020-14318
CVE-2020-14323	CVE-2020-14331	CVE-2020-14345
CVE-2020-14346	CVE-2020-14347	CVE-2020-14351
CVE-2020-14355	CVE-2020-14360	CVE-2020-14361
CVE-2020-14362	CVE-2020-14363	CVE-2020-14364
CVE-2020-14372	CVE-2020-14385	CVE-2020-14779

CVEs Fixed in Release 7.9.0.0		
CVE-2020-14781	CVE-2020-14782	CVE-2020-14792
CVE-2020-14796	CVE-2020-14797	CVE-2020-14803
CVE-2020-15436	CVE-2020-15705	CVE-2020-15706
CVE-2020-15707	CVE-2020-15862	CVE-2020-15999
CVE-2020-16092	CVE-2020-17507	CVE-2020-24394
CVE-2020-25211	CVE-2020-25212	CVE-2020-25632
CVE-2020-25637	CVE-2020-25643	CVE-2020-25645
CVE-2020-25647	CVE-2020-25648	CVE-2020-25656
CVE-2020-25684	CVE-2020-25685	CVE-2020-25686
CVE-2020-25692	CVE-2020-25694	CVE-2020-25695
CVE-2020-25705	CVE-2020-25712	CVE-2020-27749
CVE-2020-27779	CVE-2020-28374	CVE-2020-29573
CVE-2020-29599	CVE-2020-29661	CVE-2020-35513
CVE-2021-2144	CVE-2021-2163	CVE-2021-3156
CVE-2021-20225	CVE-2021-20233	CVE-2021-20265
CVE-2021-20305	CVE-2021-25215	CVE-2021-27219
CVE-2021-27363	CVE-2021-27364	CVE-2021-27365
CVE-2021-27803		

CVEs Fixed in everRun Release 7.8.0.0

The following table lists the fixed CVEs (click drop-down icon, if present)

CVEs Fixed in Release 7.8.0.0		
CVE-2015-2716	CVE-2015-8035	CVE-2015-9289
CVE-2016-5131	CVE-2017-6519	CVE-2017-11166
CVE-2017-12805	CVE-2017-12806	CVE-2017-15412
CVE-2017-15710	CVE-2017-17807	CVE-2017-18251
CVE-2017-18252	CVE-2017-18254	CVE-2017-18258
CVE-2017-18271	CVE-2017-18273	CVE-2017-18595
CVE-2017-1000476	CVE-2018-1116	CVE-2018-1301
CVE-2018-4180	CVE-2018-4181	CVE-2018-4300
CVE-2018-4700	CVE-2018-5712	CVE-2018-5745
CVE-2018-7191	CVE-2018-7418	CVE-2018-7584
CVE-2018-8804	CVE-2018-9133	CVE-2018-10177
CVE-2018-10360	CVE-2018-10547	CVE-2018-10804
CVE-2018-10805	CVE-2018-11362	CVE-2018-11439
CVE-2018-11656	CVE-2018-12599	CVE-2018-12600
CVE-2018-13139	CVE-2018-13153	CVE-2018-14340
CVE-2018-14341	CVE-2018-14368	CVE-2018-14404
CVE-2018-14434	CVE-2018-14435	CVE-2018-14436
CVE-2018-14437	CVE-2018-14567	CVE-2018-15518
CVE-2018-15587	CVE-2018-15607	CVE-2018-16057

CVEs Fixed in Release 7.8.0.0		
CVE-2018-16328	CVE-2018-16749	CVE-2018-16750
CVE-2018-17199	CVE-2018-18066	CVE-2018-18544
CVE-2018-18751	CVE-2018-19622	CVE-22018-19869
CVE-2018-19870	CVE-2018-19871	CVE-2018-19872
CVE-2018-19873	CVE-2018-19985	CVE-2018-20169
CVE-2018-20467	CVE-2018-20852	CVE-2018-21009
CVE-2019-2737	CVE-2019-2739	CVE-2019-2740
CVE-2019-2805	CVE-2019-3820	CVE-2019-3880
CVE-2019-3890	CVE-2019-3901	CVE-2019-5436
CVE-2019-6465	CVE-2019-6477	CVE-2019-7175
CVE-2019-7397	CVE-2019-7398	CVE-2019-9024
CVE-2019-9503	CVE-2019-9924	CVE-2019-9956
CVE-2019-9959	CVE-2019-10131	CVE-2019-10197
CVE-2019-10207	CVE-2019-10218	CVE-2019-10638
CVE-2019-10639	CVE-2019-10650	CVE-2019-10871
CVE-2019-11190	CVE-2019-11459	CVE-2019-11470
CVE-2019-11472	CVE-2019-11487	CVE-2019-11597
CVE-2019-11598	CVE-2019-11884	CVE-2019-12293
CVE-2019-12382	CVE-2019-12779	CVE-2019-12974

CVEs Fixed in Release 7.8.0.0		
CVE-2019-12975	CVE-2019-12976	CVE-2019-12978
CVE-2019-12979	CVE-2019-13133	CVE-2019-13134
CVE-2019-13135	CVE-2019-13232	CVE-2019-13233
CVE-2019-13295	CVE-2019-13297	CVE-2019-13300
CVE-2019-13301	CVE-2019-13304	CVE-2019-13305
CVE-2019-13306	CVE-2019-13307	CVE-2019-13309
CVE-2019-13310	CVE-2019-13311	CVE-2019-13454
CVE-2019-13648	CVE-2019-14283	CVE-2019-14815
CVE-2019-14980	CVE-2019-14981	CVE-2019-15090
CVE-2019-15139	CVE-2019-15140	CVE-2019-15141
CVE-2019-15221	CVE-2019-15605	CVE-2019-15916
CVE-2019-16056	CVE-2019-16708	CVE-2019-16709
CVE-2019-16710	CVE-2019-16711	CVE-2019-16712
CVE-2019-16713	CVE-2019-16746	CVE-2019-16865
CVE-2019-17041	CVE-2019-17042	CVE-2019-17540
CVE-2019-17541	CVE-2019-17666	CVE-2019-18634
CVE-2019-18660	CVE-2019-19338	CVE-2019-19527
CVE-2019-19768	CVE-2019-19948	CVE-2019-19949
CVE-2020-0543	CVE-2020-0548	CVE-2020-0549

CVEs Fixed in Release 7.8.0.0		
CVE-2020-1938	CVE-2020-2754	CVE-2020-2755
CVE-2020-2756	CVE-2020-2757	CVE-2020-2773
CVE-2020-2781	CVE-2020-2800	CVE-2020-2803
CVE-2020-2805	CVE-2020-2830	CVE-2020-2922
CVE-2020-5208	CVE-2020-5260	CVE-2020-5312
CVE-2020-7039	CVE-2020-8112	CVE-2020-8597
CVE-2020-8608	CVE-2020-8616	CVE-2020-8617
CVE-2020-9484	CVE-2020-10188	CVE-2020-10531
CVE-2020-10711	CVE-2020-10757	CVE-2020-10772
CVE-2020-11008	CVE-2020-12049	CVE-2020-12351
CVE-2020-12352	CVE-2020-12653	CVE-2020-12654
CVE-2020-12662	CVE-2020-12663	CVE-2020-12888
CVE-2020-14364	CVE-2020-14556	CVE-2020-14577
CVE-2020-14578	CVE-2020-14579	CVE-2020-14583
CVE-2020-14593	CVE-2020-14621	

CVEs Fixed in everRun Release 7.7.0.0

The following table lists the fixed CVEs (click drop-down icon, if present)

CVEs Fixed in Release 7.7.0.0		
CVE-2016-3186	CVE-2016-3616	CVE-2016-10713

CVEs Fixed in Release 7.7.0.0		
CVE-2016-10739	CVE-2017-5731	CVE-2017-5732
CVE-2017-5733	CVE-2017-5734	CVE-2017-5735
CVE-2017-14503	CVE-2017-17742	CVE-2018-0495
CVE-2018-0734	CVE-2018-1050	CVE-2018-1111
CVE-2018-1122	CVE-2018-1139	CVE-2018-1312
CVE-2018-3058	CVE-2018-3063	CVE-2018-3066
CVE-2018-3081	CVE-2018-3282	CVE-2018-3613
CVE-2018-5383	CVE-2018-5407	CVE-2018-5741
CVE-2018-6790	CVE-2018-6914	CVE-2018-6952
CVE-2018-7159	CVE-2018-7409	CVE-2018-7456
CVE-2018-7485	CVE-2018-7755	CVE-2018-8087
CVE-2018-8777	CVE-2018-8778	CVE-2018-8779
CVE-2018-8780	CVE-2018-8905	CVE-2018-9363
CVE-2018-9516	CVE-2018-9517	CVE-2018-10689
CVE-2018-10779	CVE-2018-10853	CVE-2018-10858
CVE-2018-10904	CVE-2018-10907	CVE-2018-10911
CVE-2018-10913	CVE-2018-10914	CVE-2018-10923
CVE-2018-10926	CVE-2018-10927	CVE-2018-10928
CVE-2018-10929	CVE-2018-10930	CVE-2018-10963

CVEs Fixed in Release 7.7.0.0		
CVE-2018-11212	CVE-2018-11213	CVE-2018-11214
CVE-2018-11645	CVE-2018-11813	CVE-2018-12015
CVE-2018-12121	CVE-2018-12181	CVE-2018-12327
CVE-2018-12404	CVE-2018-12641	CVE-2018-12697
CVE-2018-12900	CVE-2018-13053	CVE-2018-13093
CVE-2018-13094	CVE-2018-13095	CVE-2018-13346
CVE-2018-13347	CVE-2018-14348	CVE-2018-14498
CVE-2018-14598	CVE-2018-14599	CVE-2018-14600
CVE-2018-14625	CVE-2018-14647	CVE-2018-14651
CVE-2018-14652	CVE-2018-14653	CVE-2018-14654
CVE-2018-14659	CVE-2018-14660	CVE-2018-14661
CVE-2018-14734	CVE-2018-15473	CVE-2018-15594
CVE-2018-15686	CVE-2018-15853	CVE-2018-15854
CVE-2018-15855	CVE-2018-15856	CVE-2018-15857
CVE-2018-15859	CVE-2018-15861	CVE-2018-15862
CVE-2018-15863	CVE-2018-15864	CVE-2018-16062
CVE-2018-16396	CVE-2018-16402	CVE-2018-16403
CVE-2018-16646	CVE-2018-16658	CVE-2018-16838
CVE-2018-16842	CVE-2018-16866	CVE-2018-16881

CVEs Fixed in Release 7.7.0.0		
CVE-2018-16885	CVE-2018-16888	CVE-2018-17100
CVE-2018-17101	CVE-2018-17336	CVE-2018-18074
CVE-2018-18281	CVE-2018-18310	CVE-2018-18384
CVE-2018-18520	CVE-2018-18521	CVE-2018-18557
CVE-2018-18661	CVE-2018-18897	CVE-2018-19058
CVE-2018-19059	CVE-2018-19060	CVE-2018-19149
CVE-2018-19519	CVE-2018-19788	CVE-2018-20060
CVE-2018-20481	CVE-2018-20650	CVE-2018-20662
CVE-2018-20856	CVE-2018-20969	CVE-2018-1000073
CVE-2018-1000074	CVE-2018-1000075	CVE-2018-1000076
CVE-2018-1000077	CVE-2018-1000078	CVE-2018-1000079
CVE-2018-1000132	CVE-2018-1000876	CVE-2018-1000877
CVE-2018-1000878	CVE-2019-0154	CVE-2019-0155
CVE-2019-0160	CVE-2019-0161	CVE-2019-0217
CVE-2019-0220	CVE-2019-1125	CVE-2019-1387
CVE-2019-1559	CVE-2019-2503	CVE-2019-2529
CVE-2019-2614	CVE-2019-2627	CVE-2019-2945
CVE-2019-2949	CVE-2019-2962	CVE-2019-2964
CVE-2019-2973	CVE-2019-2975	CVE-2019-2978

CVEs Fixed in Release 7.7.0.0		
CVE-2019-2981	CVE-2019-2983	CVE-2019-2987
CVE-2019-2988	CVE-2019-2989	CVE-2019-2992
CVE-2019-2999	CVE-2019-3459	CVE-2019-3460
CVE-2019-3811	CVE-2019-3827	CVE-2019-3840
CVE-2019-3846	CVE-2019-3858	CVE-2019-3861
CVE-2019-3880	CVE-2019-3882	CVE-2019-3900
CVE-2019-5010	CVE-2019-5489	CVE-2019-6470
CVE-2019-7149	CVE-2019-7150	CVE-2019-7222
CVE-2019-7310	CVE-2019-7664	CVE-2019-7665
CVE-2019-9200	CVE-2019-9500	CVE-2019-9506
CVE-2019-9631	CVE-2019-9740	CVE-2019-9824
CVE-2019-9947	CVE-2019-9948	CVE-2019-10086
CVE-2019-10126	CVE-2019-10216	CVE-2019-11043
CVE-2019-11135	CVE-2019-11236	CVE-2019-11599
CVE-2019-11729	CVE-2019-11745	CVE-2019-11810
CVE-2019-11833	CVE-2019-12155	CVE-2019-13616
CVE-2019-13638	CVE-2019-13734	CVE-2019-14287
CVE-2019-14378	CVE-2019-14744	CVE-2019-14811
CVE-2019-14812	CVE-2019-14813	CVE-2019-14816

CVEs Fixed in Release 7.7.0.0		
CVE-2019-14817	CVE-2019-14821	CVE-2019-14835
CVE-2019-14869	CVE-2019-14895	CVE-2019-14898
CVE-2019-14901	CVE-2019-14906	CVE-2019-15239
CVE-2019-17133	CVE-2019-18397	CVE-2019-18408
CVE-2019-1000019	CVE-2019-1000020	CVE-2019-1010238
CVE-2020-2583	CVE-2020-2590	CVE-2020-2593
CVE-2020-2601	CVE-2020-2604	CVE-2020-2654
CVE-2020-2659		

CVEs Fixed in everRun Release 7.6.1.0

The following table lists the fixed CVEs (click drop-down icon, if present)

CVEs Fixed in Release 7.6.1.0		
CVE-2015-8830	CVE-2015-9262	CVE-2016-4913
CVE-2016-9396	CVE-2017-0861	CVE-2017-3735
CVE-2017-10661	CVE-2017-16997	CVE-2017-17805
CVE-2017-18198	CVE-2017-18199	CVE-2017-18201
CVE-2017-18208	CVE-2017-18232	CVE-2017-18267
CVE-2017-18344	CVE-2017-18360	CVE-2017-1000050
CVE-2018-0494	CVE-2018-0495	CVE-2018-0732
CVE-2018-0737	CVE-2018-0739	CVE-2018-1050

CVEs Fixed in Release 7.6.1.0		
CVE-2018-1060	CVE-2018-1061	CVE-2018-1092
CVE-2018-1094	CVE-2018-1113	CVE-2018-1118
CVE-2018-1120	CVE-2018-1130	CVE-2018-1139
CVE-2018-1304	CVE-2018-1305	CVE-2018-5344
CVE-2018-5391	CVE-2018-5407	CVE-2018-5729
CVE-2018-5730	CVE-2018-5742	CVE-2018-5743
CVE-2018-5803	CVE-2018-5848	CVE-2018-6485
CVE-2018-6764	CVE-2018-7208	CVE-2018-7568
CVE-2018-7569	CVE-2018-7642	CVE-2018-7643
CVE-2018-7740	CVE-2018-7757	CVE-2018-8014
CVE-2018-8034	CVE-2018-8781	CVE-2018-8945
CVE-2018-9568	CVE-2018-10322	CVE-2018-10372
CVE-2018-10373	CVE-2018-10534	CVE-2018-10535
CVE-2018-10733	CVE-2018-10767	CVE-2018-10768
CVE-2018-10844	CVE-2018-10845	CVE-2018-10846
CVE-2018-10852	CVE-2018-10858	CVE-2018-10878
CVE-2018-10879	CVE-2018-10881	CVE-2018-10883
CVE-2018-10902	CVE-2018-10906	CVE-2018-10911
CVE-2018-10940	CVE-2018-11236	CVE-2018-11237

CVEs Fixed in Release 7.6.1.0		
CVE-2018-11784	CVE-2018-12126	CVE-2018-12127
CVE-2018-12130	CVE-2018-12180	CVE-2018-12910
CVE-2018-13033	CVE-2018-13405	CVE-2018-13988
CVE-2018-14526	CVE-2018-14618	CVE-2018-14633
CVE-2018-14646	CVE-2018-14665	CVE-2018-15688
CVE-2018-15908	CVE-2018-15909	CVE-2018-15911
CVE-2018-16395	CVE-2018-16511	CVE-2018-16539
CVE-2018-16540	CVE-2018-16541	CVE-2018-16802
CVE-2018-16863	CVE-2018-16864	CVE-2018-16865
CVE-2018-16871	CVE-2018-16884	CVE-2018-17183
CVE-2018-17456	CVE-2018-17961	CVE-2018-17972
CVE-2018-18073	CVE-2018-18284	CVE-2018-18311
CVE-2018-18397	CVE-2018-18445	CVE-2018-18559
CVE-2018-18690	CVE-2018-19134	CVE-2018-19409
CVE-2018-19475	CVE-2018-19476	CVE-2018-19477
CVE-2018-1000007	CVE-2018-1000026	CVE-2018-1000120
CVE-2018-1000121	CVE-2018-1000122	CVE-2018-1000301
CVE-2019-2422	CVE-2019-2602	CVE-2019-2684
CVE-2019-2698	CVE-2019-2745	CVE-2019-2762

CVEs Fixed in Release 7.6.1.0		
CVE-2019-2769	CVE-2019-2786	CVE-2019-2816
CVE-2019-2842	CVE-2019-3813	CVE-2019-3815
CVE-2019-3835	CVE-2019-3838	CVE-2019-3839
CVE-2019-3855	CVE-2019-3856	CVE-2019-3857
CVE-2019-3862	CVE-2019-3863	CVE-2019-5953
CVE-2019-6116	CVE-2019-6133	CVE-2019-6454
CVE-2019-6778	CVE-2019-6974	CVE-2019-7221
CVE-2019-8322	CVE-2019-8323	CVE-2019-8324
CVE-2019-8325	CVE-2019-9636	CVE-2019-10132
CVE-2019-10160	CVE-2019-10161	CVE-2019-10166
CVE-2019-10167	CVE-2019-10168	CVE-2019-11085
CVE-2019-11091	CVE-2019-11477	CVE-2019-11478
CVE-2019-11479	CVE-2019-11811	CVE-2019-12735

CVEs Fixed in everRun Release 7.6.0.0

The following table lists the fixed CVEs (click drop-down icon, if present)

CVEs Fixed in Release 7.6.0.0		
CVE-2016-2183	CVE-2017-3636	CVE-2017-3641
CVE-2017-3651	CVE-2017-3653	CVE-2017-10268
CVE-2017-10378	CVE-2017-10379	CVE-2017-10384

CVEs Fixed in Release 7.6.0.0		
CVE-2017-11600	CVE-2017-13215	CVE-2018-1336
CVE-2018-2562	CVE-2018-2622	CVE-2018-2640
CVE-2018-2665	CVE-2018-2668	CVE-2018-2755
CVE-2018-2761	CVE-2018-2767	CVE-2018-2771
CVE-2018-2781	CVE-2018-2813	CVE-2018-2817
CVE-2018-2819	CVE-2018-2952	CVE-2018-3133
CVE-2018-3136	CVE-2018-3139	CVE-2018-3149
CVE-2018-3169	CVE-2018-3180	CVE-2018-3183
CVE-2018-3214	CVE-2018-3620	CVE-2018-3639
CVE-2018-3646	CVE-2018-3665	CVE-2018-3693
CVE-2018-5390	CVE-2018-5740	CVE-2018-7550
CVE-2018-7566	CVE-2018-8088	CVE-2018-10194
CVE-2018-10675	CVE-2018-10873	CVE-2018-10897
CVE-2018-10915	CVE-2018-11235	CVE-2018-11806
CVE-2018-12020	CVE-2018-12384	CVE-2018-14634
CVE-2018-15910	CVE-2018-16509	CVE-2018-16542
CVE-2018-1002200		

REST API

everRun systems support the Representational State Transfer (REST) application programming interface (API). With access to the REST API methods, you can collect system information, including physical

machine (PM) and virtual machine (VM) properties, statistics, and alerts. You can also execute some administrative tasks including creating VMs, powering VMs on and off, putting a PM into maintenance mode, and powering the system off. You can use these HTTP methods to create your own customized Web-based applications to monitor and manage the everRun system.

The system provides the REST API and interactive REST API documentation through the OpenAPI (also known as Swagger™) Specification. You can access the documentation directly from the local everRun system.

To begin exploring the REST API on everRun systems, see the following sections:

- Accessing the REST API documentation
- Logging In for REST API Testing
- Extending a Login Session with the Refresh Token
- Logging Out of the REST API
- List of REST API Methods

Prerequisites:

- The REST API uses port 8888 for HTTPS communications. If needed, ask your network administrator to open this port in your company's firewall to allow REST API activity.
- Optionally, to secure REST API activity, you can install a custom SSL certificate into the CA bundle of the everRun system. For information about using the `certificate_installer` script to install a custom SSL certificate, see To install a custom certificate into the CA bundle.

Accessing the REST API documentation

In a web browser, open the following URL, which includes the hostname or IP address of your everRun system:

`https://`*`hostname_or_IP address`*`:8888/restapi/swagger-ui/index.html`

The documentation describes the API methods and syntax, and also allows you to try the methods on your own system.

Logging In for REST API Testing

1. Open the REST API documentation page.
2. Under **Authentication API**, click the **/v2/login** method to expand it, and click **Try it out**.
3. The system displays the syntax for the REST API **/v2/login** method. Under **Request Body**, update the text to enter the username and password for a user account on the everRun system:

   ```
   {
   "username": "type_username_here",
   "password": "type_password_here"
   }
   ```

   Log in to an account that has the appropriate level of privileges for the REST API methods that you want to use.
4. Below the **Request Body**, click **Execute** and view the **Server Response**. If the method is successful, the response contains two REST API session authentication tokens (`token` and `refreshToken`). If the method is unsuccessful, the response might contain a message similar to `Unauthorized - Invalid username/password`.
5. Upon a successful login, copy the authorization `token` string to your management PC running the everRun Availability Console. Copy only the text between quotes to the clipboard to use in a subsequent step.

   ```
   {
   "token":
   "eyJhbGciOiJIUzUxMiJ9dgkl;jageXNUaW1lTWlsbGkiOjE2OTQzMjUzNTgwODMkilegjk",
   "refreshToken":
   "eyJhbGciOiJIUzUxMiJ9.eyJzeXNUaW1lTWlsbGkiOjE2OTkyMzQ2NDI4MTksInN1YiI6I"
   }
   ```
6. At the top of the REST API documentation page, on the right side, click **Authorize**.
7. Under **bearerAuth**, in the **Value** field, paste the authorization `token` string, and then click **Authorize**. If your login is successful, the dialog box displays a **Logout** button.
8. Click **Close** or **X** to close the dialog box and return to the REST API documentation page.
9. Proceed with your REST API testing.

Extending a Login Session with the Refresh Token

After logging in as described in Logging In for REST API Testing, the login session remains active for 45 minutes. To extend the login session, use the **/v2/refreshToken** method to reset the session timeout to 90 minutes.

1. From the original **/v2/login** response of your current login session, copy the `refreshToken` string to your clipboard. Copy only the text between quotes.

   ```
   {
   "token":
   "eyJhbGciOiJIUzUxMiJ9dgkl;jageXNUaW1lTWlsbGkiOjE2OTQzMjUzNTgwODMkilegjk",
   "refreshToken":
   "eyJhbGciOiJIUzUxMiJ9.eyJzeXNUaW1lTWlsbGkiOjE2OTkyMzQ2NDI4MTksInN1YiI6I"
   }
   ```

2. Under **Authentication API**, click the **/v2/refreshToken** method to expand it, and click **Try it out**.
3. The system displays the syntax for the REST API **/v2/refreshToken** method. Replace the `string` text between the quote marks with the `refreshToken` string you copied in the previous step:

   ```
   {
   "refreshToken": "string"
   }
   ```

4. Below the **Request Body**, click **Execute** and view the **Server Response**.
5. From the response of the **/v2/refreshToken** method, copy the new authorization `token` string to your clipboard. Copy only the text between quotes.

   ```
   {
   "token":
   "eyJhbGciOiJIUzUxMiJ9.eyJzeXNUaW1lTWlsbGkiOjE2OTkyMzQ2NDI4MTksInN1Yi",
   }
   ```

6. At the top of the REST API documentation page, on the right side, click **Authorize**, and then click **Logout**.
7. Under **bearerAuth**, in the **Value** field, paste the new authorization `token` string, and then click **Authorize**. If your login is successful, the dialog box displays a **Logout** button.
8. Click **Close** or **X** to close the dialog box and return to the REST API documentation page.
9. Proceed with your REST API testing.

Logging Out of the REST API

At the top of the REST API documentation page, click **Authorize** and then click **Logout**.

List of REST API Methods

The following methods are available:

VirtualMachine Api

POST /v2/virtualmachines/vmshutdown

POST /v2/virtualmachines/vmpoweron

POST /v2/virtualmachines/vmpoweroff

POST /v2/virtualmachines/vmforceboot

POST /v2/virtualmachines/vmcreate

POST /v2/virtualmachines/vmcopy

GET /v2/resource/virtualmachines

GET /v2/resource/virtualmachines/*vmname*

DELETE /v2/virtualmachines/vmdelete

UpgradePath Api

GET /v2/upgradepath

Resource Api

GET /v2/resource

Summary Api

GET /v2/summary

Logout Api

PUT /v2/logout

Default Api

GET /v2/

VirtualCD Api

POST /v2/vcd/savedescription

POST /v2/vcd/edit

POST /v2/vcd/detach

POST /v2/vcd/create

POST /v2/vcd/attach

GET /v2/vcd/vcdlist

GET /v2/vcd/download/*vcdName*

DELETE /v2/vcd/remove

Preference Api

GET /v2/pref

PhysicalMachine Api

POST /v2/physicalmachines/workon

POST /v2/physicalmachines/workoff

POST /v2/physicalmachines/systemshutdown

POST /v2/physicalmachines/nodeshutdown

GET /v2/physicalmachines

GET /v2/physicalmachines/*pmname*

Authentication Api

POST /v2/refreshtoken

POST /v2/login

PreferenceSystem Api

GET /v2/pref/system

GET /v2/pref/system/swupdate

GET /v2/pref/system/quorumservers

GET /v2/pref/system/productlicense

GET /v2/pref/system/ownerinfo

GET /v2/pref/system/mailserver

GET /v2/pref/system/ipconfiguration

GET /v2/pref/system/dateandtime

Alert Api

GET /v2/alert

GET /v2/alert/history

PreferenceRemote Api

GET /v2/pref/remotesupport

SupportLogs Api

GET /v2/supportlogs

GET /v2/auditlogs

System Api

GET /v2/system/node

GET /v2/system/cluster

GET /v2/system/listusbs

VMImport Api

POST /v2/import/listfiles

POST /v2/import/importvm

PreferenceAdmin Api

GET /v2/pref/admintools

GET /v2/pref/admintools/secureconnection

PreferenceNotification Api

GET /v2/pref/notification

12

Chapter 12: Security

To learn about additional configuration settings that you can implement to provide the highest level of security for a everRun system, see Security Hardening.

For additional information about security, see the following topics:

- Fixed CVEs
- Managing IPtables
- Configuring Secure Connections
- Configuring Users and Groups
- Configuring Active Directory
- Audit Logs Page

Security Hardening

Although Stratus everRun software provides a secure out-of-box experience, you can implement additional configuration settings as described below to provide the highest level of security.

Security is often a balance between protection and ease of use. everRun software is shipped with a set of default settings that balance these factors. For a more secure posture, follow the guidelines below, and continue to evaluate the security of the system throughout its life cycle, from the planning and configuration to operation and decommissioning.

The information below provides security hardening guidance based on Version 7.1 of *CIS Controls*, which are hardening recommendations developed by the Center for Internet Security (CIS), a community-driven nonprofit that leads and is recognized for best practices for securing IT systems and data. *CIS Benchmarks*

are also used to validate and create a baseline for a secure product. A list of CIS Controls is included below in Best Practices and Standards of Standards Organizations.

The information below also provides hardening guidance based on industrial control systems cyber security standard ISA/IEC 62443, which was originally created by the International Society of Automation (ISA) and continues to be developed by the International Electrotechnical Commission (IEC). ISA/IEC 62443-4-2 has differing levels of security based on the sensitivity of data or intended threat actor adversary, and by implementing the recommendations and applying mitigating controls assists in achieving compliance for the required security level. A summary of ISA/IEC 62443-4-2 requirements is included below in Best Practices and Standards of Standards Organizations.

This help topic contains the following sections:

- Security Guidelines
- Advanced Security Guidelines
- Best Practices and Standards of Standards Organizations

Security Guidelines

The following sections describe security guidelines for everRun systems.



Note: Stratus has tested and supports the following guidelines. Any other update or modification not explicitly approved by Stratus could affect the normal operation of the system.
If you have any questions about these guidelines, and the system is covered by a service agreement, contact your authorized Penguin Solutions service representative for assistance. For information, see the **everRun Support** page at https://www.stratus.com/services-support/customer-support/?tab=everrun

While implementing the security hardening guidelines, consider the following:

- The security guidelines refer to administrative tasks performed in the everRun Availability Console and in the host operating system. The everRun Availability Console is a browser-based interface that allows you to manage and monitor most aspects of the everRun system from a remote management computer (see the everRun Availability Console). The host operating system runs on each node of the system. You can access the command line of the host operating system locally at the PM's physical console or remotely by using a secure shell (SSH) client (see Accessing the Host Operating System).

- Prior to making any configuration changes, record the current settings so that you can restore them, if necessary. Also, record any modifications that you are making in case the information is needed for troubleshooting.
- When changing the default system settings, particularly in the host operating system, you must make the changes on both nodes to prevent inconsistencies that could affect the normal operation of the system. Similarly, when changing the `root` password and other user account settings for the host operating system, you must do so on both nodes. The guidelines below indicate when these changes are needed.
- When you upgrade the system software or replace a node in the system, not all modifications for system hardening may be carried over. Similarly, some settings are shared across nodes, so shared resources could have conflicts. Therefore, after completing these procedures, you should verify that each node in the system has the correct settings and that the system is working properly.
- In some cases, the security guidelines directly reference Knowledge Base articles (for example, KB*nnnnnn*) with more information about configuring everRun systems and the everRun software. You can access the **Penguin Solutions Customer Service Portal** and its Knowledge Base by using your existing portal credentials, or by creating a new user account, as described in Accessing Knowledge Base Articles.

Ports and Protocols

Any administrator making networking or communication changes to the system should be knowledgeable about the ports or protocols used by everRun. For details, see KB0012595.

Network Segmentation

Connect the everRun system only to networks with trusted devices, or to networks where devices require explicit permissions to communicate with each other. For more information on network segmentation, see the NIST special publications 800-125B and 800-39. For information about which Ethernet networks are available on everRun systems, see Network Architecture.

IP Tables/Firewall

Enable IP tables packet filtering for the system, and block all ports that are not used in normal operation. Malicious actors can leverage a potential security vulnerability on an unused interface as a backdoor. Limit the exposure by enabling IP tables for unused ports.

For details on how to implement IP tables, see Managing IPtables.

Notes:

- The ICMP protocol is used for pinging within the everRun system. If you set IP tables to drop ICMP traffic, the fault tolerance or failover support will not work properly.
- The SSH protocol is used for connecting to the host operating system. If you set IP tables to block SSH traffic, system administrators will be unable to access the host operating system.

User Account Creation

Create individual user accounts for each user authorized to access the system, and consider each user's role in the usage of the device. Maintaining individual user accounts also permits auditability or nonrepudiation, that by log review it can be determined which user accessed the device or made configuration changes.

For details on how to configure user settings, see Configuring Users and Groups.

Notes:

- You cannot delete the default **admin** account, although you should change its name and password by editing the account settings.
- You must specify an email address for each user account, including **admin**, to enable the forgot password feature. Also, you must enable the mail server, as described in Configuring the Mail Server; otherwise, the system cannot send password reset emails.
- If a user account is no longer needed or not actively being used, either remove or disable the user account to prevent any possibility of inappropriate use.
- Monitor login attempts to prevent brute-force attacks.

Password Creation

You must change the default passwords for the system.

The everRun Availability Console prompts you for a new **admin** password upon deployment. The password policy of the everRun Availability Console requires that your password meets the following conditions:

- Its minimum length is 8 characters.
- It must contain both upper- and lower-case characters.
- It cannot be the username.

Also change the `root` password for the host operating system on both nodes as soon as possible. To change the password, issue the `passwd` command on each node. For details, see Accessing the Host Operating System.



Note: When you change the `root` password for the host operating system, ensure that you remember the password, because the only way to recover a lost `root` password is to replace or reinstall the nodes.

For more information about controlling the quality of passwords in the host operating system, see Advanced Security Guidelines.

Least Privilege

Limit each user's access to features applicable to their position or role.

Implementing least privilege prevents a non-privileged user from accessing services above their role.

For details on how to configure roles that define the privileges for each user, see Configuring Users and Groups.

Active Directory

Active Directory integration presents a single point for centralized authentication and authorization. With Active Directory, you can create group policies for password complexity that are enforced based on your local security policy.

For details on how to add a everRun system to an Active Directory domain, see Configuring Active Directory.

BIOS

The BIOS is an important aspect of the operation of the system; therefore, you should implement password protection for the BIOS to protect it from malicious manipulation.

You can set a password for the BIOS by accessing the system BIOS setup utility while the system is booting.

Note: When you change the BIOS password, ensure that you remember the password. If you lose the BIOS password, typically the only way to restore access is to reset the BIOS and manually restore the configuration settings.

Ports

Disable any system ports that are not required for normal operation.

Time Synchronization

Synchronization of time is important, as it provides a centralized reference point to ensure that operation and security processes work within the same time frame. Time referencing allows for confidence in the time of check and time of use when updating applications and ensuring that keys and certificates are still valid based on the time and date.

When you log on to a everRun system for the first time, enable the Network Time Protocol (NTP) service to automatically set the system clock . Configure NTP to reference a known and trusted NTP server. For details, see Configuring Date and Time.

Note: Use only the everRun Availability Console to properly configure the NTP settings; do not manually configure them in the host operating system.

Secure Connections

By default, the everRun Availability Console is configured to support only secure connections with the HTTPS protocol.

Enabling HTTPS on the everRun system prevents common web security attacks to provide a level of confidentiality for each web session. HTTPS encrypts web session traffic, provides data integrity, and increases the overall security of the web traffic.

When HTTPS is enabled, it supports only TLSv1.2, which is currently the strongest encryption suite recommended. Ciphers include:

TLSv1.2:
ciphers:
TLS_DHE_RSA_WITH_AES_128_CBC_SHA (dh 4096) - A
TLS_DHE_RSA_WITH_AES_128_CBC_SHA256 (dh 4096) - A
TLS_DHE_RSA_WITH_AES_128_GCM_SHA256 (dh 4096) - A

TLS_DHE_RSA_WITH_AES_256_CBC_SHA (dh 4096) - A
TLS_DHE_RSA_WITH_AES_256_CBC_SHA256 (dh 4096) - A
TLS_DHE_RSA_WITH_AES_256_GCM_SHA384 (dh 4096) - A
TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA (dh 4096) - A
TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA (dh 4096) - A
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA (secp256r1) - A
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256 (secp256r1) - A
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 (secp256r1) - A
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA (secp256r1) - A
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384 (secp256r1) - A
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 (secp256r1) - A

Also enable secure, encrypted connections when using a mail server or other types of server software. Beginning with everRun 7.9.1.0, the system requires that the mail server you configure for e-Alerts and password resets supports the TLS v1.2 protocol. If your mail server does not support TLS 1.2, then no outgoing emails will be sent, even if they are configured in the everRun Availability Console. For information about configuring and enabling an encrypted connection for the mail server on a everRun system, see Configuring the Mail Server.

Updating SSL Certificate

The everRun system comes with a self-signed SSL certificate, but this may be updated to any purchased or supplied certificate. Changing the SSL certificate allows the root of trust to be updated to the customer specification. For details, see KB0013477.

SNMP Configurations

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).

For security reasons, SNMP is disabled by default on everRun systems. In everRun Release 7.9 or higher, the SNMP process is also stopped in the console operating system on each node. For additional security, you can also disable all SNMP connections by adding rules to IPtables (see Managing IPtables) to block UDP ports 162, 161 and 199 and TCP ports 162 and 199.



Note: For security reasons, if you need to enable SNMP, you should disable SNMP v1 and v2, and enable only version 3 by using the SNMP **Restricted** configuration. For details, see Configuring SNMP Settings.

Backups

Backups are important to have in case a security event occurs; a unit can be returned to a known good state for continuous operation. Any backups taken should be stored in a secure location.

To back up a running VM to a set of XML and raw hard disk data files, see Backing Up a Running Virtual Machine. To recover the identical VM with the same SMBIOS UUID, system serial number, and MAC addresses as the original VM, see Recovering a Virtual Machine from a Backup. Alternatively, you can export a stopped VM to a set of OVF and VHD files and restore them as described in Exporting a Virtual Machine and Replacing/Restoring a Virtual Machine from an OVF File.

On redundant, dual-node everRun systems, each node also serves as a backup for the other node. If a node fails, you can replace a node in a system that is currently licensed, and the system automatically restores the node with an exact copy of the everRun software and the virtual machines from the running node.

SplitSite Configuration

A SplitSite configuration connects two physical machines at 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, it is very important that you use the quorum service because a SplitSite configuration exposes the A-Link networks to other potential failure scenarios.

For details, see Creating a SplitSite Configuration.

Auditing

Implement auditing by a local policy to regularly collect and manage logs of events needed to detect, understand, and recover from a cyber attack.

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.)

Log information contains:

- 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 by using `snmptable` (for details, see Obtaining System Information with snmptable).

Use logs for continuous monitoring of the everRun system. To ensure prompt service in the event of a service call, also enable support notifications and periodic reporting for your system to keep Stratus informed about your system's health. For details, see Configuring Remote Support Settings.

Login Banner Notice

Configure the Login Banner Notice to include important notifications to everRun Availability Console users. For details, see Configuring the Login Banner.

Upgrades

Upgrade everRun on a regular basis to prevent security vulnerabilities from being exploited due to out-of-date components. Refer to your local security policies for information about frequency and methods.

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

The **Upgrade Kits** page in the everRun Availability Console allows 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 Using an Upgrade Kit. For information about creating a USB medium, see Creating a USB Medium with System Software.

Physical Security

Install each everRun system in a secure location to prevent malicious users from accessing the nodes.

Secure each location with an auditable system to identify which personnel entered the area to identify malicious users.

Physical security is an important addition to tamper detection and alerting for any device, including everRun nodes.

Advanced Security Guidelines

The following sections describe advanced security guidelines for everRun systems.

Password Quality Recommendations

When setting passwords, recommendations include:

- Setting a minimum password length of at least 8 characters, of which three out of four of the following characteristics are required: one upper-case letter, one lower-case letter, one number, and one special character.
- Requiring users to reset passwords on a regular basis, such as every 30, 60 or 90 days. You can also forbid the reuse of passwords for a variable amount of password updating history.

To manually update password quality settings in the host operating system



Note: Apply the password quality settings on both nodes in the system.

1. Log on to the host operating system, as described in Accessing the Host Operating System.
2. Open the `/etc/pam.d/common-password` file with a text editor.
3. Modify the `pam_pwquality.so` module with the appropriate settings. For example, use settings similar to the following:

```
password requisite pam_pwquality.so try_first_pass local_
users_only retry=3 authtok_type= minlen=8 lcredit=-1
ucredit=-1 dcredit=-1 ocredit=-1 enforce_for_root
```

The previous example sets the following values:

`minlen=8` sets the minimum password length to 8 characters.

`lcredit=-1` sets the minimum number of lower-case letters in a password to one.

`ucredit=-1` sets the minimum number of upper-case letters in a password to one.

`dcredit=-1` sets the minimum number of digits in a password to one.

`ocredit=-1` sets the minimum number of other symbols such as @, #, ! $ % in a password to one.

`enforce_for_root` ensures that even if the `root` user is setting the password, the complexity policies should be enforced.

4. To restrict the password history, add or modify the `pam_pwhistory.so` module with the appropriate settings. For example, using settings similar to the following:

```
password requisite pam_pwhistory.so debug use_authtok
remember=10 retry=3
```

5. Save the `/etc/pam.d/common-password` file.

For more information about user management and password policies in the host operating system, see the Ubuntu documentation:

https://ubuntu.com/server/docs/security-users

Concurrent User Management

Continually monitor the audit logs to view which users have logged on to the machine and if they are still active.

Identify the users that are currently operating the system to legitimize and audit their usage.

Antivirus

Continually perform a network-based analysis for antivirus or malware detection.

Your network-based intrusion detection system supplements the everRun capability to support verification of the intended operation of security functions. The detection system should search for anomalous network traffic and require investigation to validate any malicious intent.

SSH Access Restrictions

Several `/etc/ssh/sshd_config` parameters limit which users and groups can access the system by SSH. If none of the following parameters are present in the file, edit the file to set one or more of them to limit access:

`AllowUsers`

The `AllowUsers` parameter gives the system administrator the option of allowing specific users to use SSH to access the system. The list consists of space separated usernames. This parameter does not recognize numeric user IDs. To restrict user access further by permitting only the allowed users to log in from a host, the entry can be specified in the form of user@host.

`AllowGroups`

The `AllowGroups` parameter gives the system administrator the option of allowing specific groups of users to use SSH to access the system. The list consists of space separated group names. This parameter does not recognize numeric group IDs.

`DenyUsers`

The `DenyUsers` parameter gives the system administrator the option of denying specific users from using SSH to access the system. The list consists of space separated usernames. This parameter does not recognize numeric user IDs. If a system administrator wants to restrict user access further by specifically denying a user's access from a host, the entry can be specified in the form of user@host.

`DenyGroups`

The `DenyGroups` parameter gives the system administrator the option of denying specific groups of users from using SSH to access the system. The list consists of space separated group names. This parameter does not recognize numeric user IDs.

Restricting which users can remotely access the system using SSH will help ensure that only authorized users access the system.

`MaxAuthTries`

The `MaxAuthTries` parameter specifies the maximum number of authentication attempts permitted per connection. When the login failure count reaches half the number, error messages will be written to the `syslog` file detailing the login failure.

Setting the `MaxAuthTries` parameter to a low number will minimize the risk of successful brute force attacks to the SSH server. While the recommended setting is 4, set the number based on site policy. For example:

```
MaxAuthTries 4
```

`IgnoreRhosts`

The `IgnoreRhosts` parameter specifies that `.rhosts` and `.shosts` files will not be used in `RhostsRSAAuthentication` or `HostbasedAuthentication`.

Setting this parameter forces users to enter a password when authenticating with SSH. For example:

```
IgnoreRhosts yes
```

`HostbasedAuthentication`

The `HostbasedAuthentication` parameter specifies if authentication is allowed through trusted hosts by using `.rhosts` or `/etc/hosts.equiv` with successful public key client host authentication. This option applies only to SSH Protocol Version 2.

Even though the `.rhosts` files are ineffective if support is disabled in `/etc/pam.conf`, disabling the ability to use `.rhosts` files in SSH provides an additional layer of protection. For example:

```
HostbasedAuthentication no
```

For more information about `sshd_config` parameters, see the `sshd_config(5)` manual page.

Best Practices and Standards of Standards Organizations

The information in this topic is based on the following best practices and standards.

CIS Controls version 7.1

CIS controls is a prioritized set of best practices created to stop the most pervasive and dangerous threats of today. It was developed by leading security experts from around the world and is refined and validated every year. Further details may be found on the CIS website: https://www.cisecurity.org.

The CIS controls are:

Basic

1. Inventory and Control of Hardware Assets
2. Inventory and Control of Software Assets
3. Continuous Vulnerability Management
4. Controlled Use of Administrative Privileges
5. Secure Configuration for Hardware and Software on Mobile Devices, Laptops, Workstations and Servers
6. Maintenance, Monitoring and Analysis of Audit Logs

Foundational

7. Email and Web Browser Protections
8. Malware Defenses
9. Limitation and Control of Network Ports, Protocols and Services
10. Data Recovery Capabilities
11. Secure Configuration for Network Devices, such as Firewalls, Routers and Switches
12. Boundary Defense
13. Data Protection
14. Controlled Access Based on the Need to Know
15. Wireless Access Control
16. Account Monitoring and Control

Organizational

17. Implement a Security Awareness and Training Program
18. Application Software Security
19. Incident Response and Management
20. Penetration Tests and Red Team Exercises

ISA/IEC 62443-4-2

ISA/IEC 62443-4-2 details technical component requirements (CRs) associated with seven foundational requirements (FRs) for meeting control system capability security levels. Further details may be found on the IEC website: https://www.iec.ch/

The foundational requirements are:

1. Identification and authentication control (IAC)
2. Use control (UC)
3. System integrity (SI)
4. Data confidentiality (DC)
5. Restricted data flow (RDF)
6. Timely response to events (TRE)
7. Resource availability (RA)

1. Identification and authentication control (IAC)

Identification of users is used in conjunction with authorization mechanisms to implement access control for a component. Verifying the identity of users requesting access is necessary to protect against unauthorized users from gaining access to the component. Authorization is from access control lists for different users that log in and authenticate with passwords into the everRun system.

2. Use control (UC)

Once the user is identified and authenticated, the component must restrict the allowed actions to authorized use of the component. The everRun system has defined roles that implement the concept of least privilege. Creating multiple users with varying levels of access control also defines the authorized use of the component.

3. System integrity (SI)

The integrity of the device should not be compromised, both the software and the physical components in operational and non-operational states. For example, the everRun system validates the digital signatures of software components prior to an upgrade. Ensuring system integrity is important to protect against the unauthorized manipulation or modification of data or system.

4. Data confidentiality (DC)

The purpose is to ensure the confidentiality of information on communication channels and in data stored in repositories to protect against unauthorized disclosure. The everRun system has HTTPS with TLS v1.2 for web communication, as well as SSH and SMTP with encryption, ensuring that information is protected from malicious persons.

5. Restricted data flow (RDF)

Restricted data flow is the segmentation of the control system through zones and conduits to limit the unnecessary flow of data. The everRun network architecture supports the routing and switching as determined by the configuration of networking for the management of information flow as determined by the installed system engineer. Leveraging the networking capabilities of the everRun system allows for network segmentation to limit data flow.

6. Timely response to events (TRE)

Although a system may begin operation in a secure state, vulnerabilities and security events can occur. The everRun system has a Product Security Incident Response (PSIR) team to react to security incidents and report findings while solving issues in a timely manner. The everRun system has alert logs that can be used to notify the appropriate channels for configuration changes that may indicate a security incident. The logs contain enough information for forensics, and these e-alert notifications are emailed.

7. Resource availability (RA)

The aim of this control is to ensure that the component is resilient against various types of denial of service events. The high availability of the everRun system is the foundation of an “always on” state. It is imperative that industrial control systems maintain a high availability state as there potentially are life safety impacts to systems. With a built-in virtualization and availability layer, automated data protection, and application recovery, everRun significantly reduces the dependence on IT for virtualized computing at the edge. Its self-protecting and self-monitoring features help reduce unplanned downtime and ensure the continuous availability of business-critical industrial applications.

13

Chapter 13: SNMP

Simple Network Management Protocol (SNMP) is a standard protocol for receiving alarms, sending traps, and monitoring system status. SNMP draws upon system-defining information that is stored in hierarchically configured management information bases (MIBs).

To configure an everRun system to use SNMP, see Configuring SNMP Settings.

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.penguinsolutions.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
```

The command output displays the following:

Field	Description
everRunAlertIndex	The alert number.
everRunAlertSeverity	The alert severity (see everRunAlertSeverityNum for numerical value). Values are: clear  informational  minor  major  serious  critical 
everRunAlertType	The type of alert. Examples are: • node_singleSystemDisk • Node Maintenance • The Unit is not well balanced
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 *yyyy-mm-dd hh:mm:ss*, where *yyyy* is year, *mm* is month, *dd* is date, *hh* is hour, *mm* is minute, and *ss* is second (for example, 2017-11-03 23:49:45).
everRunAlertCallHomeSent	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
`everRunAlertSNMPTrapSent`	If `true`, SNMP trap was sent; if `false`, it was not sent
`everRunAlertInformation`	Information about the alert. Examples are: • `Node node1 is in maintenance` • `node0 has a single system disk: Policy assumes this disk is redundant - if not, please add another internal disk` • `BUSINESS network net_728 is reporting a degraded link condition` • `The unit is not well load balanced`
`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 *yyyy-mm-dd hh:mm:ss*, where *yyyy* is year, *mm* is month, *dd* is date, *hh* is hour, *mm* is minute, and *ss* is second (for example, `2017-11-03 23:49:45`).
`everRunAuditUsername`	The name of the user that caused the audit to be generated (for example, `audit` or `admin`).
`everRunAuditOriginatingHost`	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
```

The command output displays the following:

Field	Description
everRunNodeIndex	A number (typically 1 or 2) to indicate the node whose information is displayed.
everRunNodeId	The host ID of the node (for example, host:o34).
everRunNodeDisplayName	The node name, node0 or node1.
everRunNodeIsPrimary	If true, the node is primary. If false, the node is secondary.
everRunNodeStateNum	Node state is: 0 Normal () 1 Warning () 2 Busy () 3 Broken () 4 Maintenance ()
everRunNodeActivityNum	Node activity is: 0 Imaging 1 Booting 2 Running 3 Stopping 4 Rebooting 5 Powered off 6 Failed 7 Firmware updating 8 Lost 9 Exiled 10 Unreachable

Field	Description
	11 Proto (initializing) 12 Evacuating

To display VM information

To display VM information, issue the following command:

```
snmptable -v2c -m+/usr/smd/STRATUS-EVERRUN-MIB.txt -c public
localhost everRunVMTable
```

The command output displays the following:

Field	Description
everRunVMIndex	An incrementing number (1, 2, etc.) to indicate the VM whose information is displayed.
everRunVMId	The VM ID (for example, vm:o1467).
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
```

The command output displays the following:

Field	Description
everRunVolumeIndex	An incrementing number (1, 2, etc.) to indicate the volume whose information is displayed.
everRunVolumeId	The volume ID (for example, volume:o588).
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 ()