Quorum®

onQ Administrator's Guide

Version 3.8 SP1 HF1

Quorum®
2890 Zanker Road, Suite 102
San Jose, CA 95134 USA
Toll Free: +1-877-99-Quorum
Phone: +1-408-708-4500
Fax: +1-408-708-4503

Copyright © 2008–2015 QuorumLabs, Inc.

ALL RIGHTS RESERVED. No part of this publication may be reproduced in any form, by photocopy, microfilm, retrieval system, or by any other means now known or hereafter invented without the prior written permission of QuorumLabs, Inc.

The information in this publication has been carefully checked and is believed to be accurate. However, QuorumLabs, Inc. assumes no responsibility for inaccuracies. QuorumLabs, Inc. retains the right to make changes to this publication at any time without prior notice. QuorumLabs, Inc. does not assume any liability arising from the application or use of this publication or the product(s) described herein.

TRADEMARKS

All trademarks or registered trademarks used in this publication are the property of their respective owners.

Table of Contents

Preface 1

Online Help 1

Chapter Organization 1

1. Product Overview 3

1.1 About onQ 3

1.2 Terminology 6

1.3 How does onQ work? 7

1.4 Inside onQ Appliance 8

1.5 About Hybrid Cloud 9

1.6 About onQ Flex 10

1.7 Product Support 10

2. Requirements 11

2.1 Configuration Guidelines & Requirements 11

2.2 Network and Firewall Requirements 13

2.2.1 LAN/Network Communications 13

2.2.2 onQ Appliance-to-PN Communications 14

2.2.3 HA-to-DR Communications 14

2.2.4 WAN-DR Site Communications 16

2.3 Host-based Firewall and Anti-Virus Requirements 16

2.4 Operational Considerations 18

2.5 Identify interdependencies 18

3. Deployment 19

3.1 Installation Overview 19

3.2 (Workflow) Install onQ Appliances 21

3.2.1 Step 1: Prepare to install 23

3.2.2 Step 2: Configure the onQ Appliances 23

3.2.3 Step 3: (Hybrid Cloud Only) Create VPN connection to Hybrid Cloud 25

3.2.4 Step 4: Connect the Appliance to your network 35
3.2.5 Step 5: Change hypervisor's NTP server 36
3.2.6 Step 6: Enroll protected nodes 36
3.2.7 Step 7: Configure RN startup dependencies 36
3.2.8 Step 8: Start automated protection process 36
3.3 Initial Backup 37
3.4 Initialize the DR Repository 38

4. The Basics 41
4.1 Log on to Appliance's user interface 42
4.2 Establish a VPN connection to Hybrid Cloud 43
4.3 Log off from Appliance's user interface 44
4.4 Drop-down menu 44
4.5 Reboot onQ Manager 46
4.6 Reboot Appliance 47
4.7 Shut down or Restart Appliance 47
4.8 About onQ Proxy 48
4.9 Modify Appliance's network settings 49
4.10 Configure Appliance's hypervisor settings 50
4.11 Unlock UI 52
4.12 (Agent-based PNs) Restart the onQ Service 54
4.13 Modify protected nodes 55
4.14 Delete protected nodes 56
4.15 Protected Node Parameters 58

5. Enrollment in non-cluster Environment 63
5.1 (Agent-based Windows PNs) Enroll protected nodes 64
5.2 (Agent-based Linux PNs) Enroll protected nodes 72
5.3 (Agent-less Linux/Windows PNs) Enroll protected nodes 105
5.4 (Agent-based PNs) Add protected nodes manually 136

6. Enrollment in Windows Cluster Services Environment 139

6.1 Windows Cluster Services Overview 139
6.2 (Step 1) Prepare the cluster nodes 140
6.3 (Step 2) Install onQ Service on cluster 143
6.4 (Step 3) Configure cluster for the onQ Service 147
6.5 (Step 4) Configure the onQ Appliance for cluster environment 153
6.6 (Step 5) Enroll the cluster on the onQ Appliance 154
6.7 Work with RNs in cluster environment 156

7. **Updates** 163
7.1 (Start Here) Upgrade Support and Requirements 163
7.2 Update Appliance software 165
7.3 (Agent-based Windows PNs) Update node software 167
7.4 (Agent-based Linux PNs) Update node software 169
7.5 (Agent-less Windows/Linux PNs) Update PN Proxy software 172
7.6 Get software packages 173
7.7 Delete software packages 176
7.8 (Agent-based Linux PNs) Uninstall node software packages 177
7.9 Install onQ Appliance license 177
7.10 View onQ software versions 179
7.11 Specify onQ Central Proxy 180

8. **Protection** 181
8.1 Start node protection 182
8.2 Stop node protection 182
8.3 Restart node protection 183
8.4 Modify recovery node creation policy 184
8.5 Recreate recovery nodes 186
8.6 Disable protection of nodes 189
8.7 Enable protection of nodes 190
8.8 Create on-demand RNs 191
8.9 Delete on-demand RNs 192

8.10 Cancel RN updates 193
8.11 Disable replication globally 194
8.12 Disable replication for individual nodes 195
8.13 Enable replication globally 197
8.14 Enable replication for individual nodes 198
8.15 Manage hosts 198
8.16 Resize protected node's vdisk 200
8.17 Enable and disable DR Mirroring 201
8.18 Enroll and Disenroll onQ Archive Vault 204

9. Backup and Restore 205
9.1 Backup and Recovery Workflow 206
9.2 Edit RN services list 208
9.3 Edit backup exclude list 210
9.4 Edit backup include list 213
9.5 Initiate immediate backups 214
9.6 Export or import Repository 215
9.7 Schedule backups 217
9.8 Change Backup Retention Policy 218
9.9 Delete snapshots 219
9.10 Stop in-progress backups 220
9.11 Retrieve objects from Repository 222
9.11.1 Perform WSR Restore 222
9.11.2 Perform File-level Restore 225
9.12 Adjust bandwidth throttling 235
9.13 Limit resource utilization on PNs 236
9.14 Run custom backup scripts 237
9.15 (Agent-based PNs) Methods of Performing Incremental Backups 241
9.16 Perform full rescan on PN 243
9.17 Enable onQ to back up shared volumes 244

10. Security and Communications 245
10.1 Test HA-to-DR Link 246
10.2 Set up trust relationships 247
10.3 Create Custom Networks for RNs 252
10.4 Assign RNs to Networks 255
10.5 Modify TCP Ports for HA-DR Communications 259
10.6 Synchronize system time 259
10.7 Secure connection to PNs 260

11. Disaster Recovery and DR Preparedness 263
11.1 Configure startup dependencies 264
11.2 (Workflow) Fail over HA Appliance to DR Appliance 268
11.3 (Workflow) Fail back DR to HA Appliance 270
11.4 (Workflow) Fail over a PN to an RN 272
11.5 (Workflow) Fail back an RN to a PN 274
11.6 Change Appliance's role 275
11.7 Start recovery nodes on HA Appliance 277
11.8 Start recovery nodes on DR Appliance or DR Mirror 285
11.9 Stop recovery nodes 288
11.10 Back up and restore Oracle 11g database on Linux 290
11.11 Back up and restore Oracle 10g+ database on Windows 300
11.12 Cancel or override RN startup process 310
11.13 Launch recovery node's console 311
11.14 Test Mode 312
11.15 Configure automatic testing of RNs 313
11.16 Suspend and resume RN tests 313
11.17 Restore PNs Using QUARK 315

12. Monitoring 337
12.1 Monitor multiple onQ Appliances 338
12.2 Generate on-demand reports 341

12.3 Monitor protected nodes 358
12.4 Monitor recovery nodes 365
12.5 Monitor backups 369
12.6 Backup Activity Report 370
12.7 Monitor DR Appliance 373
12.8 Monitor onQ Archive Vault enrollment 379
12.9 Monitor disk space and memory usage 379
12.10About Alerts 381
12.11Modify e-mail alert settings 383
12.12Retrieve logs 384

13. Troubleshooting 387
13.1 Installation Problems 388
13.2 (Agent-based PNs) Connection Problems 389
13.3 (Agent-less PNs) Connection and Backup Problems 391
13.4 Recovery Node Problems 394
13.5 Self-Test and RN Boot Problems 395
13.6 (Agent-based PNs) Protected Node Upgrade Problems 401
13.7 Helpful System Information 401
13.8 Generate Support Log 402
13.9 Backup Alerts 403
13.9.1 A0801 403
13.9.2 A0802 408
13.9.3 A0803 408
13.10Snapshot Alerts 409
13.10.1A0401 409
13.10.2A0402 409
13.10.3A0403 410
13.10.4A0404 410
13.11PN Status Alerts 411
13.11.1A0704 411
13.11.2A0710 412

13.12License Expiration and Upgrade Alerts 412
13.12.1A0008 413
13.12.2A0009 413
13.12.3A0014 413
13.12.4A0015 414
13.12.5A0017 414
13.13Reboot Alerts 414
13.13.1A0010 415
13.13.2A0011 415
13.13.3A0012 415
13.13.4A0013 415
13.13.5A0018 416
13.14Protection Alerts 416
13.14.1A0001 416
13.14.2A0003 417
13.15Recovery Node and PN Disk Space Alerts 417
13.15.1A0901 417
13.15.2A0902 418
13.15.3A0903 418
13.15.4A0904 419
13.15.5A0906 419
13.16Repository Space 419
13.16.1A0101 419
13.16.2A0102 420
13.16.3A0103 420
13.16.4A0104 421
13.16.5A0303 421
13.16.6A0304 421
13.17onQ Disk Space Alerts 422
13.17.1A0201 422
13.17.2A0301 422
13.17.3A0302 423
13.18Connection Alerts 423
13.18.1A0705 423

13.18.2A0706 424
13.18.3A0707 424
13.18.4A0708 424
13.19Self-Test Alerts 425
13.19.1A1001 425
13.19.2A1002 425
13.19.3A1003 426
13.20Change iDRAC network configuration 426

14. User Management 427
14.1 Add users 427
14.2 Delete users 428
14.3 Change user passwords 429

15. Example Configurations 431
15.1 Local onQ Appliance Configurations 432
15.1.1 Local: Example of HA Appliance 432
15.1.2 Local: Example of DR Appliance 433
15.1.3 Local: Example of DR Mirror 433
15.2 Remote onQ Appliance Configurations 433
15.2.1 Remote: Example of HA Appliance with Hybrid Cloud 435
15.2.2 Remote: Example of HA Appliance with Remote DR Appliance 437
15.2.3 Remote: Example of DR Appliance with DR Mirror 438
15.2.4 Remote: Example of DR Mirror 440

16. Example Logs 441
16.1 Example Event Log 442
16.2 Example Event DB Log 443
16.3 Example Expired Snapshot Log 444
16.4 Example FLR Activity Log 444
16.5 Example HA -> DR Transfer Log 446
16.6 Example Manager Debug Log 447
16.7 Example PN Configuration Log 448

16.8 Example Self Test Log 449
16.9 Example Upgrade Log 450
16.10Example WSR Activity Log 451

17. Tech Notes 1
17.1 Install kernel-xen RPM package 1

18. Glossary 5

Preface

This guide is intended for System Administrators who need to in install, configure, manage, and monitor their onQ Appliances. The guide can be used by both first-time and experienced users.

Online Help

Online Help for onQ and Archive Vault products are available through their respective portals. Online documentation is available at:

Help	Where to go...
onQ	http://www.quorum.net/help/3.8SP1HF1/onQ/index.html
AV	http://www.quorum.net/help/3.8SP1HF1/Vault/index.html

Chapter Organization

The chapters and appendices of this manual are as follows:

- Chapter 1, *Product Overview*
- Chapter 2, *Requirements*
- Chapter 3, *Deployment*
- Chapter 4, *The Basics*
- Chapter 5, *Enrollment in non-cluster Environment*
- Chapter 6, *Enrollment in Windows Cluster Services Environment*
- Chapter 7, *Updates*
- Chapter 8, *Protection*
- Chapter 9, *Backup and Restore*
- Chapter 10, *Security and Communications*
- Chapter 11, *Disaster Recovery and DR Preparedness*
- Chapter 12, *Monitoring*
- Chapter 13, *Troubleshooting*

- Chapter 14, *User Management*
- Chapter 15, *Example Configurations*
- Chapter 16, *Example Logs*
- Chapter 17, *Tech Notes*
- Chapter 18, *Glossary*.

1

Product Overview

- About onQ
- Terminology
- How does onQ work?
- Inside onQ Appliance
- About Hybrid Cloud
- About onQ Flex
- Product Support

1.1 About onQ

Businesses today depend upon the continuous availability of information systems and business applications. Any downtime is disruptive and expensive.

An effective business continuity plan is essential for any business that wants to remain viable and competitive even when systems fail or business is interrupted.

Typical business continuity solutions require a variety of technologies provided by different vendors, and involve physical duplication of production resources at remote locations. This approach works but is very expensive.

The advent of new technologies such as virtualization, replication, storage, and de-duplication have reduced the cost of business continuity for some businesses, but the overall cost and degree of expertise required has still kept these newer approaches beyond the reach of most small and medium businesses.

Quorum has created an affordable, all-in-one business continuity solution that includes all the technologies required to protect and restore physical and virtual machines. This solution is delivered on a single physical device called,

simply, onQ.

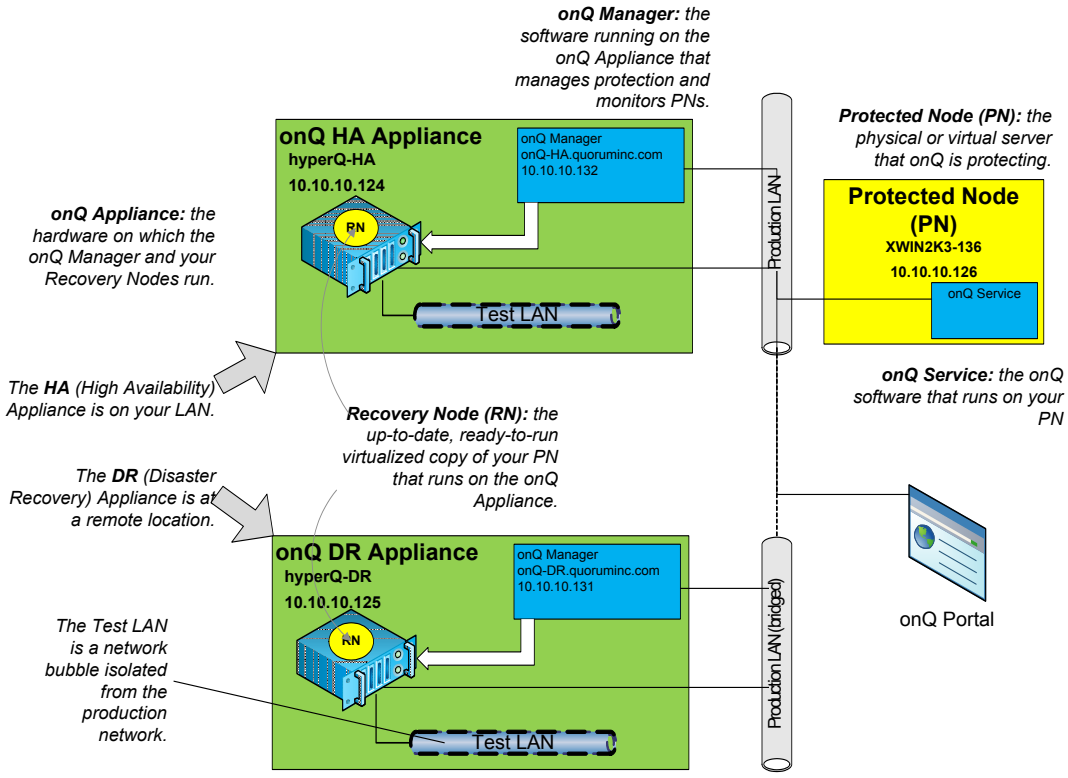
onQ is based on Quorum's feature-rich management and automation platform and comes complete with all the technologies necessary to provide high-availability and disaster-recovery protection for any Windows-based workload. With its intuitive user interface onQ is simple to install and manage.

onQ automates high availability (HA), disaster recovery (DR) and business continuity (BC) in a turnkey Appliance by integrating all of the following features and benefits:

- **Full system imaging** of all Operating System, Application and Data files to allow bare metal restoration of your servers.
- **Incremental snapshot archiving**: crash-consistent updates of all files, even if they are open, allows efficient archiving as frequently as every 15 minutes.
- **Deduplication**: only the changes on your server are saved in the archive, and only if they are different than any of the “chunks” of data already stored.
- **Replication**: your entire archive is replicated to a second remote onQ Appliance over an encrypted, compressed, bandwidth-controlled link.
- **Virtualization**: onQ maintains an up-to-date virtual copy of your servers ready-to-run right on the onQ Appliance.
- **Recovery Infrastructure**: you don’t need to have spare, standby servers—the onQ Appliance has the resources to run the ready-to-run virtual copies of your servers.
- **Monitoring, Reports and Alerts**: onQ Manager monitors all activities and lets you know when things go wrong—and when they go right.
- **Single Web-based Interface**: accessible from anywhere, the onQ Portal lets you manage all operations without complexity.

1.2 Terminology

The figure below shows some of the terminology we use when discussing onQ and its components. For a complete list of terms, go to Glossary.

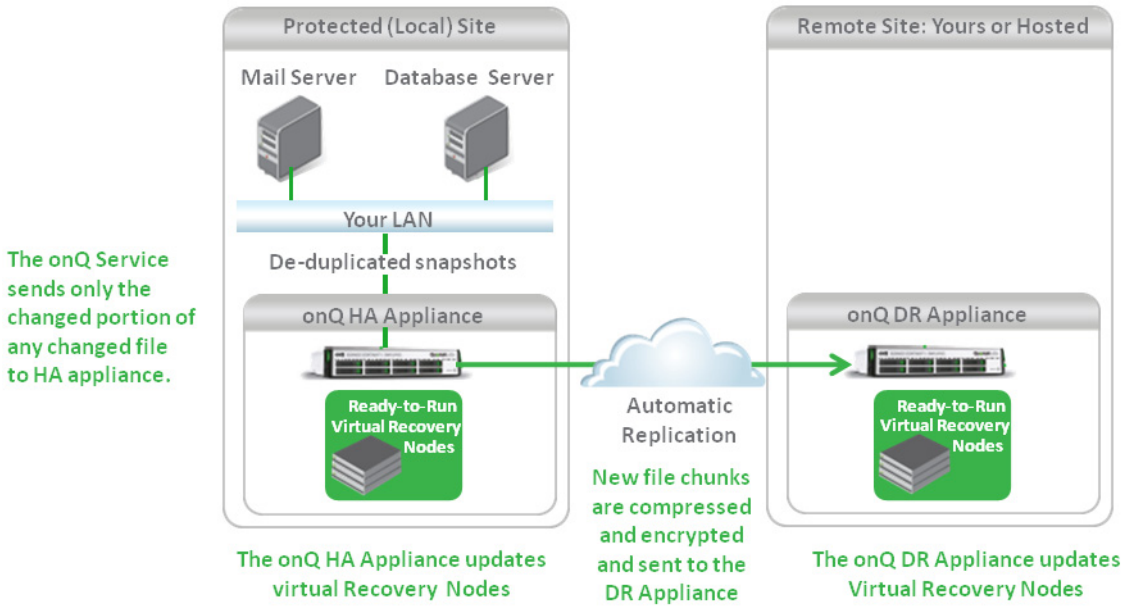


1.3 How does onQ work?

Basically, onQ works like this:

- **onQ** is a fully integrated turnkey Appliance solution. One onQ Appliance is deployed for high-availability (HA) protection. Most businesses add a second onQ Appliance at another location for disaster recovery (DR) or use Quorum's hybrid cloud offering to host their DR.
- Each onQ Appliance hosts the **onQ Manager**, Quorum's core software.
- Your HA Appliance is connected to the same network as your protected nodes. The DR Appliance is typically connected to the HA Appliance (not the protected nodes) through a VPN over a WAN.
- One of the following:
  - (*agent-based PNs*) **onQ Service** software is installed on each protected node (PN). The **onQ Manager** communicates with this software to determine operational status and to schedule snapshots. For more information go to (Agent-based PNs) Restart the onQ Service and (Agent-based Linux PNs) Enroll protected nodes.
  - (*agent-less PNs*) **Proxy hosts** are installed on the ESX/ESXi host. The **onQ Manager** communicates with these proxy hosts to determine operational status and to schedule snapshots. For more information, go to (Agent-less Linux/Windows PNs) Enroll protected nodes.
- You manage **onQ Manager** through an intuitive, browser-based, user interface called the **onQ Portal**. When you install your onQ Appliance(s), you use the UI to configure the onQ Appliance to take incremental snapshots for each protected node. You also use the UI to control and monitor the protection process.
- Once protection is started, the HA Appliance manages your backup schedule, taking deduplicated, incremental snapshots of your protected nodes at specified intervals.
  - After each snapshot is added to the repository, the onQ Manager on the HA Appliance updates the recovery node corresponding to the protected node.
  - At the same time, the DR Appliance is synchronized and independently updates its own recovery node. Thus, there are individual up-to-date recovery nodes on both the HA Appliance and the DR Appliance.

- onQ Manager alerts you by e-mail if a protected node or an incremental snapshot fails.

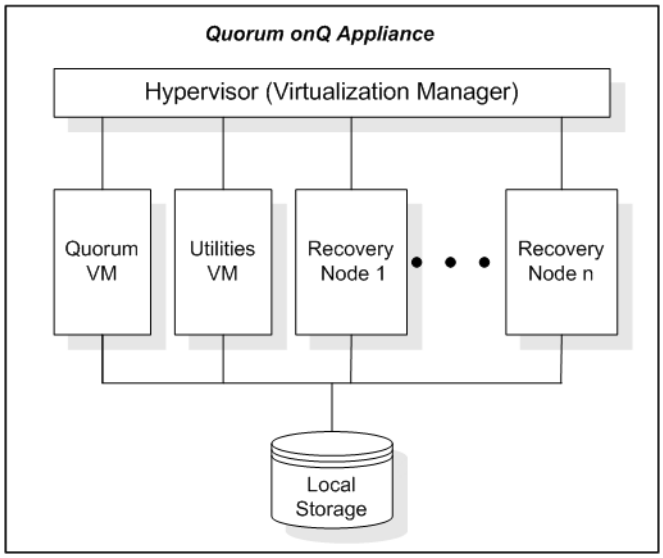


1.4 Inside onQ Appliance

Each Quorum onQ Appliance is a server class x86-based computer running on a "bare metal hypervisor" for efficient automated virtualization.

The onQ Appliance typically contains several terabytes of redundant internal disk storage. This internal storage is used to efficiently store deduplicated snapshots of your protected nodes and to host the ready-to-run recovery nodes.

onQ is organized as a set of virtual machines. Each onQ Appliance runs several virtual machines (VMs), including one VM that hosts the onQ Manager and multiple, automatically-maintained VMs that are your ready-to-run recovery nodes built from snapshots of your protected nodes.



1.5 About Hybrid Cloud

Hybrid Cloud is a DR Appliance as a Service (DRaaS). The Deployment of a Hybrid Cloud-enabled HA Appliance is virtually identical to an on-premise HA Appliance. The only difference is that you must also add information to the Hybrid Cloud-enabled HA Appliance so that it can VPN to your DRaaS. This secure connection enables the HA Appliance to send snapshots to your DRaaS.

In the event of an HA site disaster, Quorum can ship you the cloud DR Appliance, replacing it with a new DR instance. This onsite DR Appliance assumes the HA role.

The Quorum Data Center that will be hosting your applications are SAS-70 Type II certified. All data transferred from the HA Appliance to the DR Appliance goes through a 128-bit AES encrypted session behind a 256-bit AES VPN tunnel that is connected directly from the HA Appliance to the DR Appliance. There is a dedicated virtual firewall isolating each individual custom virtual network. All connections to the DR Appliance or the DR RNs are via VPN to that firewall.

Moreover, Hybrid Cloud is certified PCI compliant. As the first vendor in the industry to offer cloud recovery that meets the payment card industry's (PCI) security specifications, Quorum ensures credit card data obtained by retail

establishments, and any other organization that handles credit card information, is kept secure throughout the backup and recovery process.

1.6 About onQ Flex

onQ™ Flex is a Hybrid Cloud configuration where your business-critical Recovery Nodes (RNs) are Ready-to-Run in the event of a disaster and the remaining RNs are preconfigured to Build-on-Demand (BoD), yet all data is protected. If you want to change BoD RNs to Ready-to-Run RNs, you can do so, but there's a monthly charge; to avoid the charge, you can easily upgrade.

onQ is unable to test BoD RNs because a recovery node must be available to do so; nonetheless, in the event that you desire to test an RN or you have a disaster, simply manually create the on-demand RN; there is a monthly cost associated with creating and storing on-demand RNs, so be sure to delete the on-demand RN if you no longer need them. Lastly, unlike Ready-to-Run RNs, BoD RNs do not receive RN updates.

1.7 Product Support

Quorum Support is committed to providing the best product support in the computer industry. Customer support is available 6am–6pm (PST), Monday through Friday, via telephone, fax, e-mail, and our Web site.

Quorum Support:

Telephone:	+1.408.708.4502
E-mail:	support@quorum.net
Web Site:	http://www.quorum.net/support

Sales & Marketing:

The Sales & Marketing Department can be reached at info@quorum.net.

2

Requirements

- Configuration Guidelines & Requirements
- Network and Firewall Requirements
- Host-based Firewall and Anti-Virus Requirements
- Operational Considerations
- Identify interdependencies

2.1 Configuration Guidelines & Requirements

Before you install onQ you will need to gather some information about your site. These preparation steps are usually completed by your Reseller or Quorum using Quorum Site Survey tools and guidelines.

If you are performing this function yourself, you can request the onQ Site Survey kit from Quorum Support which includes forms and tools to streamline your deployment.

- **Network Configuration:** For all servers you wish to protect, whether virtual or physical, you will need host names and IP addresses. You will also need two addresses for each of your onQ Appliances: one for the hypervisor and the other for onQ.
- **Operating System:** The Operating System on your Protected Nodes must be supported. See "Platform Support" in onQ Release Notes.
- **Data Usage:** During setup, onQ will check storage usage for each disk on each server you protect. By default, onQ assigns 1.5 times the amount of space currently in use. You can increase this if you anticipate higher growth in the next few years.
- **Memory:** onQ uses an algorithm to assign RAM based on how much is used by your server and what is available on the onQ Appliance. If you decide to change this default allocation, consider allocating the minimum memory required for each Recovery Node for adequate operation. Only

exceed 2GB per node when you know that it is necessary for that node's workload.

- **Processors:** onQ uses an algorithm to assign virtual CPUs based on how many CPU cores are on your server and how many are available on the onQ Appliance. Because there are multiple Recovery Nodes operating concurrently on each onQ Appliance and not all will have peak activity at the same time, more virtual CPUs can be allocated than are physically present. Allocating more virtual CPUs does not necessarily improve performance and can actually hurt performance if Recovery Nodes are reserving CPUs that they do not need.
- **Microsoft VSS**. Microsoft VSS must be enabled and have adequate snapshot space allocated. For more information, go to "Protected Node Parameters" on page 58.
- **Database backups and restores**. Verify that your database backup and restore scripts are performing as expected. You should not have any issues using these same scripts in an onQ environment because such operations are independent of onQ. Do not deploy onQ if you are having database backup or restore failures. Resolve these issues before you proceed with your deployment.

2.2 Network and Firewall Requirements

Ensure that the following ports are open on the network, onQ Appliance, and PNs:

- LAN/Network Communications
- onQ Appliance-to-PN Communications
- HA-to-DR Communications
- WAN-DR Site Communications

Note: Traffic to onQ Appliances should not be open to any public access, except 443 and 123.

2.2.1 LAN/Network Communications

The following port(s) are required for administrative access to the portal on each onQ Appliance and for alerts, updates, and licensing. These ports should not be open to public access:

Direction	Port	Protocol	Purpose
LAN-only Inbound	80 & 443	TCP	onQ Portal access. Needed internally on private network LAN only. No Internet access recommended.
Inbound	443	TCP	https from onQ IP needed to the Internet for licensing and alerts
Outbound	443	HTTPS	Access for alerts (httpsalerts.onqcentral.com), updates and licensing (updates.onqcentral.com).
Inbound & Outbound	123	UDP	Communication with NTP (ntp.org).

2.2.2 onQ Appliance-to-PN Communications

Generally, onQ-to-PN communications are over a LAN that is not secured. In the event that there are any filtering devices between the PN and the onQ Appliance, the following ports need to be accessible on the PN:

Direction	Port	Protocol	Purpose
Outbound	3000	TCP	Allows squirtcopy communications
Inbound	5000	TCP	For FLR processing, allows WVhds for Windows and Netcat for Linux to receive data from onQ Manager
Inbound	5990	TCP	Allows DCRM to accept connections from onQ Manager
Inbound	5990	UDP	Allows LAMD to communicate with onQ Manager
Inbound	5991	UDP	Allows LAMD to monitor the PN
Inbound	5992	UDP	Allows LAMD to send PN heartbeat

2.2.3 HA-to-DR Communications

The following ports must be accessible between the HA Appliance and the DR Appliance pairs across WAN or LAN:

Direction	Port	Protocol	Purpose
Inbound & Outbound	22	TCP	Allows secure TCP connections between onQ Appliances. May be remapped.

Direction	Port	Protocol	Purpose
Inbound & Outbound	81	TCP	Allows inter-appliance link monitoring. May be remapped.

2.2.4 WAN-DR Site Communications

The following ports must be accessible/inaccessible between your WAN and DR site:

Direction	Port	Protocol	Purpose
Inbound & Outbound	22	TCP	<DR IP> to <HA IP>
Inbound & Outbound	81	TCP	<DR IP> to <HA IP>
Inbound	80	TCP	Deny ANY PUBLIC / External / Internet to <DR IP> (HTTP protocol)
Inbound	443	TCP	(https) Allow <DR IP> outbound to Public / External / Internet
Outbound	443	SSL, SMTP, & SSH, TCP	Access for alerts (alerts.onqcentral.com), updates (updates.onqcentral.com), and licensing (lic.onqcentral.com).

2.3 Host-based Firewall and Anti-Virus Requirements

The following four files require specific attention:

- C:\Program Files\Quorum\QuorumDCRM-NODE\acd.exe
- C:\Program Files\Quorum\QuorumDCRM-NODE\lamd.exe
- C:\Program Files\Quorum\usr\squirtcopy.exe
- C:\Program Files\Quorum\usr\WVhdS.exe

For each file, do the following:

- Create local firewall and antivirus exceptions to allow inbound and outbound access for these files across all networks—local, private, and public.
- Exclude these files (and the entire `\Quorum` folder) from any local, realtime antivirus software products.
- Any host-based firewalls local to the Protected Node (PN) should either have the same ports opened, or local firewall exceptions added for these files in order for them to function in private, local, and public modes.

2.4 Operational Considerations

- Adjust bandwidth throttling, if necessary, to speed up data transfers.
- Perform initial backup on LAN, not WAN, to improve performance. Quorum suggests that you perform the initial setup, and then the initial backup, with both your HA Appliance and DR Appliance on the same LAN. An alternative technique for performing the initial backup is to use a USB drive or other intermediate storage.
- Identify interdependencies.

2.5 Identify interdependencies

Beware of interdependent resources on your Protected Nodes. For example, if your mail server depends on a Domain Controller to function properly, both machines must be protected by onQ. Similarly, if there are critical functions that depend on DNS, ensure that onQ protects your DNS server too; otherwise, those functions might fail. onQ itself generally maintains critical IP addresses in the onQ Appliance's HOSTS file.

In the event of an HA site disaster, the boot order for your recovery nodes must take into account these dependencies. During the installation workflow, you will be prompted to configure these startup dependencies using the onQ Portal.

3

Deployment

- Installation Overview
- (Workflow) Install onQ Appliances
- Initial Backup
- Initialize the DR Repository

3.1 Installation Overview

Most installations use two Appliances. One onQ Appliance is deployed to provide high availability (HA) protection while the other is deployed in a more remote location to provide disaster recovery (DR) protection.

The two onQ Appliances are architecturally identical, although they can have different physical resources available, such as processors and memory. If your installation uses two onQ Appliances, you'll make two passes through some of the installation steps—once for the HA Appliance and once for the DR Appliance. Also, during installation, you will assign the onQ Appliance its role as the HA Appliance or the DR Appliance.

Before you can use onQ you need to install the onQ Appliances on your network and enroll your Protected Nodes (PNs). You make these initial settings and manage the ongoing operations of the onQ Appliance using the onQ Portal.

Some settings, such as network parameters, are usually pre-configured at the factory and you rarely, if ever, need to change them.

Setting up the onQ Appliances requires a special user, `VARAdmin`, with privileges to perform the following tasks. `VARAdmin` is not needed for day-to-day administration and operation.

- Changing the onQ Appliance network configuration
- Setting up Security

- Setting the VARAdmin password
- Preparing the onQ Appliance for re-deployment

You will need to log on as VARAdmin in order to configure your Appliances. After you log on, you should change this password. Contact Quorum Support for the default password.

Warning: It is critical that you remember your password: Quorum Support cannot recover it and cannot perform certain support tasks without it.

3.2 (Workflow) Install onQ Appliances

Use this procedure to set up your onQ Appliances. These instructions apply to each onQ Appliance, whether it provides high availability (HA) or disaster recovery (DR) protection.

Install your onQ Appliance in the environment in which it is to be used. If you have problems during setup, go to Installation Problems.

To set up your onQ Appliances, perform the following sequence of tasks:

Task	For instructions, Go to...
1. Get an overview of the installation process.	Installation Overview
2. Collect information about your site(s)	Step 1: Prepare to install
3. If not pre-configured, configure the onQ Appliance.	Step 2: Configure the onQ Appliances
4. (Hybrid Cloud Only) Create a VPN connection to your DR Appliance so that you can access its web portal.	Step 3: (Hybrid Cloud Only) Create VPN connection to Hybrid Cloud
5. Connect the onQ Appliance to your network.	Step 4: Connect the Appliance to your network
6. Change hypervisor's NTP server	Step 5: Change hypervisor's NTP server
7. Enroll the protected nodes you want to protect.	Step 6: Enroll protected nodes

Task	For instructions, Go to...
8. Configure the boot order for your recovery nodes.	Step 7: Configure RN startup dependencies
9. Start protection on the nodes.	Step 8: Start automated protection process
10. (Optional) Install onQ Archive Vault, if you intend to archive backups.	After the hardware is up and running, refer to the AV online help for configuration instructions. For supported configurations, go to Archive Vault Configuration Support.
11. (Recommended) Test your existing backup and recovery procedures using an RN in test mode.	See also: • Back up and restore Oracle 11g database on Linux • Back up and restore Oracle 10g+ database on Windows

3.2.1 Step 1: Prepare to install

The setup process requires a supported browser. Go to Browser Support.

To prepare to install:

1. Enable JavaScript on your browser.
2. Ensure that the list of trusted websites maintained by the Internet Explorer Enhanced Security Configuration (IEESC) includes **about.blank** and the onQ Appliance IP addresses.

   IEESC is a Windows component only shipped with Windows Server.
3. Collect information about your site using the worksheets available on the Customer Support Portal.

   Read the requirements and best practices outlined in Requirements.

Next Step: If your Appliances were pre-configured prior to delivery, skip to Step 4: Connect the Appliance to your network; otherwise, proceed to Step 2: Configure the onQ Appliances.

3.2.2 Step 2: Configure the onQ Appliances

If your Appliances were pre-configured prior to delivery, skip to Step 4: Connect the Appliance to your network. Otherwise, perform the following procedure on the HA Appliance, DR Appliance, and if applicable, DR Mirror.

To configure an onQ Appliance:

Perform the following steps on all Appliances, including the HA Appliance, DR Appliance and, if applicable, the DR Mirror.

1. Use a standalone computer with an available Ethernet port and configure the IP address to **169.254.2.99** and net mask **255.255.255.0**.
2. Connect your computer's Ethernet port directly to Port 0 on the rear of the onQ Appliance.
3. Log on to the onQ Portal with user ID `VARAdmin`. Contact Quorum Support for the default password. If you are prompted to clear the EULA acceptance, click **Yes**.
4. For security purposes, change both the default `VARAdmin` and `Admin` passwords. Go to Change user passwords.

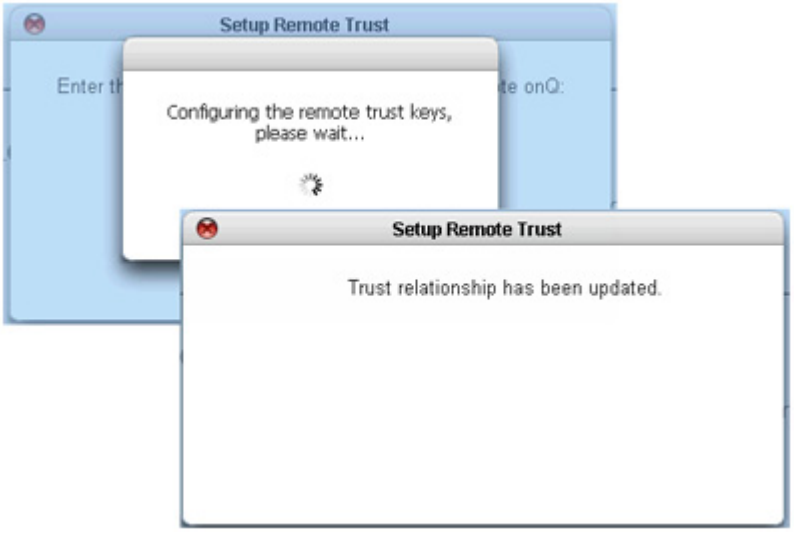
5. If Protection is on, stop protection.
6. Specify local onQ Appliance configuration:
   a. Go to **APPLIANCE CONFIG** tab > **onQ (LOCAL)** page, then **MODIFY**.
   b. Specify the parameters, including networking information and onQ Role, then **SAVE**.

      If your PNs need to go through a proxy, go to About onQ Proxy.

   For examples:
   - Local: Example of DR Appliance
   - Local: Example of DR Mirror
7. Specify remote onQ Appliance configuration:
   a. Go to **APPLIANCE CONFIG** tab > **onQ (REMOTE)** page, then **MODIFY**.
   b. Click **MODIFY** and set the parameters, including networking information, DR mirroring, and bandwidth limits, then **SAVE**.

   For examples:
   - Remote: Example of HA Appliance with Hybrid Cloud
   - Remote: Example of HA Appliance with Remote DR Appliance
   - Remote: Example of DR Appliance with DR Mirror
   - Remote: Example of DR Mirror
8. Set up the trust relationship.

You must set up the HA-to-DR link. If you have a DR Mirror, set up the DR-to-Mirror link too. For more information, go to Set up trust relationships.



a. Test the link.

9. Configure e-mail alerts.
10. Configure the hypervisor.
11. Reboot onQ.

Next Step: Step 3: (Hybrid Cloud Only) Create VPN connection to Hybrid Cloud

3.2.3 Step 3: (Hybrid Cloud Only) Create VPN connection to Hybrid Cloud

To access your Hybrid Cloud, you must do so using a VPN connection. Afterward, you can launch the web portal from your browser using the DR Appliance's LAN IP address that your Quorum Support Engineer provided you.

The PPTP VPN connection to your Hybrid Cloud is very much dependent on the Internet performance available to you and the system performance of the computer from which you are initiating the connection.

For support with connecting to the Hybrid Cloud, contact Quorum Support at

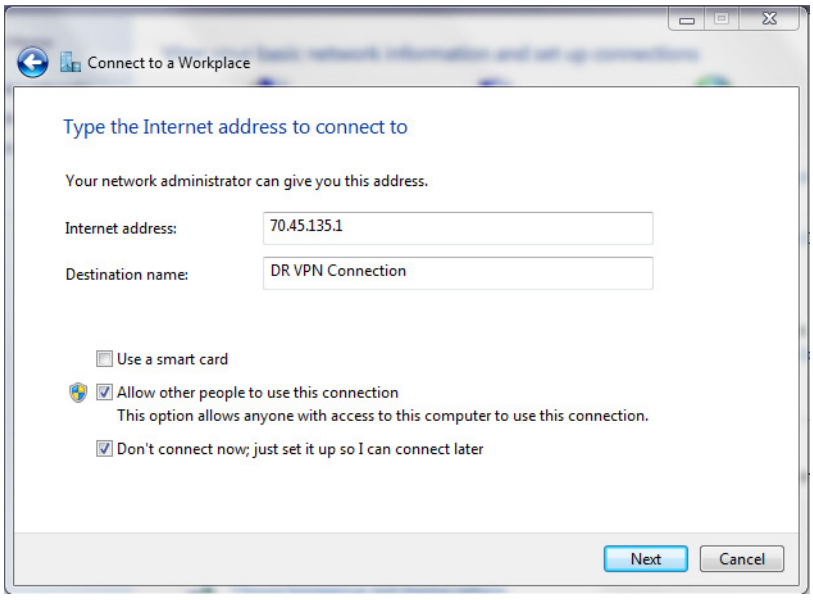
support@quorum.net.

To create a connection to Hybrid Cloud (Windows 7):

This procedure assumes that the client on which you are creating this VPN connection is running Windows 7 with XenCenter installed.

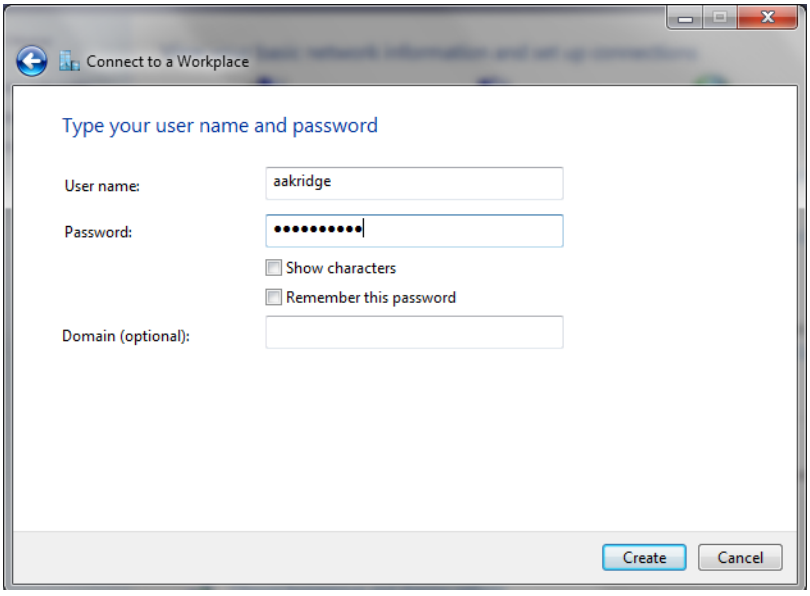
1. Retrieve the following information:
   - Public IP address and logon credentials of the DR Appliance. This information was provided by your Quorum Support Engineer.
   - LAN IP address of DR Appliance in the Quorum Data Center
2. Configure your network firewall to allow PPTP Outbound Traffic.
3. Run Software Update to ensure that your installation is running the latest updates.
4. From the Control Panel, launch the **Network and Sharing Center**, then click **Set up a new connection or network** link.
5. In the **Choose a connection option** pane, choose `Connect to a workplace`, then **Next**.
6. In the **How do you want to connect** pane, select `Use my Internet connection (VPN).`
7. Complete the following fields and options as follows, then **Next**.

Internet Address	Provide the public IP that Quorum Support provided you.
Destination Name	Type a name for the connection. For example, `DR VPN Connection`.
Allow other people to use this connection	Select this check box.
Don't Connect now; just set it up so I can connect later	Select this check box.



8. Type in your user name and password that your Quorum Support Engineer provided you and leave the domain field blank, then **Create**.

   Do not connect yet. Click **Close**.



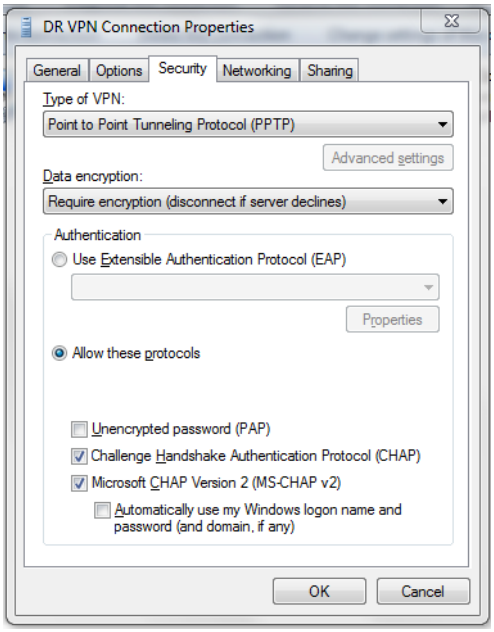
9. Configure the connection to use PPTP protocol:

   a. From **Network and Sharing Center**, click on the **Change Adapter Settings** link in the left pane.

b. Right-click on **DR VPN Connection**, then **Properties**.

c. From the **Security** tab > **Type of VPN** drop-down, choose `Port to Port Tunneling Protocol (PPTP).`

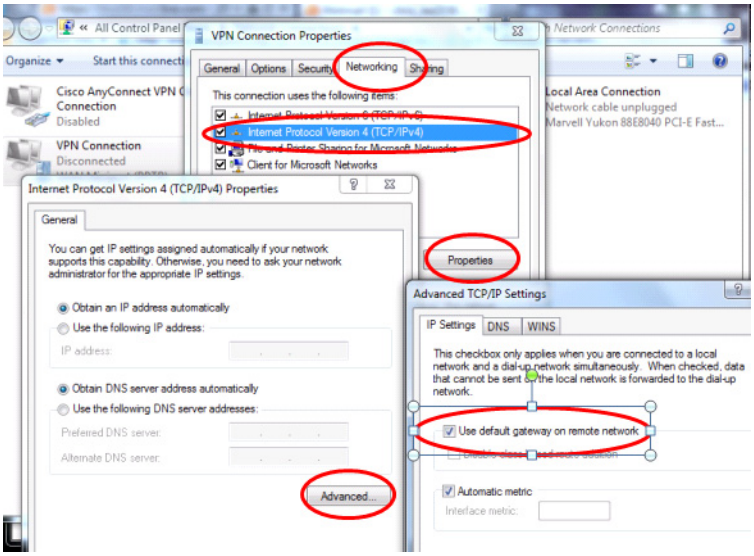
d. Accept all remaining defaults. No further settings are required, then **OK**.



10. Set remote gateway as the default gateway:

a. Right-click on **DR VPN Connection**, then **Connect**.

b. Go to the **Networking** tab > **IPv4** > **Properties** > **Advanced**. Deselect the **Use default gateway on remote network** check box.



11. Establish the VPN connection.
12. Verify that the DR Appliance is up. From a command line prompt, ping the DR Appliance's LAN IP address.
13. Log on to the DR Appliance's web portal. From a web browser, go to `https://<DR-Appliance-LAN-IP-Address>`.

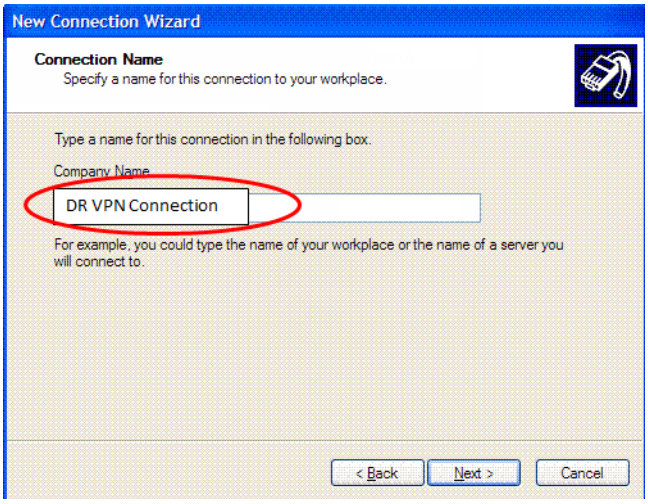
    You should now have full access to the RNs and File Level Backups on the DR Appliance.

To create a connection to Hybrid Cloud (Windows XP):

This procedure assumes that the client on which you are creating this VPN connection is running Windows XP with XenCenter installed.

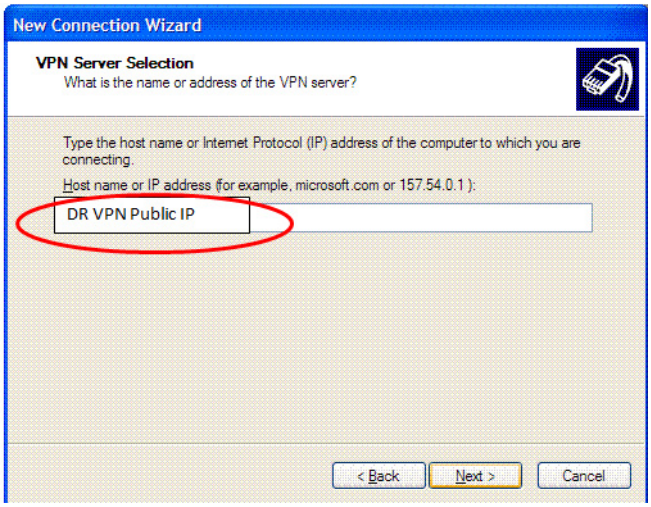
1. Retrieve the following information:
   - Public IP address and logon credentials of the DR Appliance. This information was provided by your Quorum Support Engineer.
   - LAN IP address of DR Appliance in the Quorum Data Center.
2. Configure your network firewall to allow PPTP Outbound Traffic.
3. Run Software Update to ensure that your installation is running the latest updates.

4. From the Control Panel, launch the **Network and Sharing Center**, then click **Create a new connection** link.
5. In the New Connection wizard, click **Next**.
6. Select **Connect to the network at my workplace** radio button, then click **Next**.
7. Select **Virtual Private Network connection** radio button, then **Next**.
8. In the **Company Name** text box, type `DR VPN Connection` as the connection name, then **Next**.



9. If prompted select the **Do not dial the initial connection** radio button, then **Next**.

10. In the **Host name or IP Address**... text box, type the DR Appliance's Public IP address provided by your Quorum Support Engineer, then **Next**.



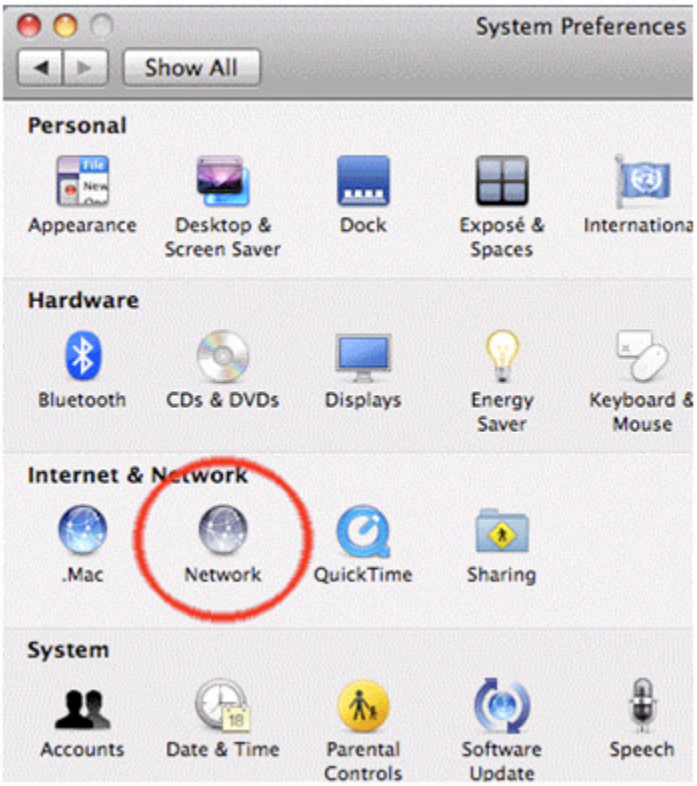
11. If prompted select either **Anyone's use** or **My use only**, then Next, then **Finish**.
12. Modify the connection's settings:
    a. Click on the **Properties** button.
    b. Click on the **Networking** tab.
    c. Highlight **Internet Protocol (TCP/IP)**, then **Properties**.
    d. Click on the **Advanced** button.
    e. Clear the **Use default gateway on remote network** check box, then **OK** until you return to the logon window, accepting all previously defined settings.
13. Establish the VPN connection.
14. Verify that the DR Appliance is up. From a command line prompt, ping the DR Appliance's LAN IP address.
15. Log on to the DR Appliance's web portal. From a web browser, go to `https://<DR-Appliance-LAN-IP-Address>`.

You should now have full access to the RNs and File Level Backups on the DR Appliance.

To create a connection to Hybrid Cloud (Mac OS X):

This procedure assumes that the client on which you are creating this VPN connection is running Mac .OS X.

1. Retrieve the following information:
   - Public IP address and logon credentials of the DR Appliance. This information was provided by your Quorum Support Engineer.
   - LAN IP address of DR Appliance in the Quorum Data Center
2. Configure your network firewall to allow PPTP Outbound Traffic.
3. Run Software netUpdate to ensure that your installation is running the latest updates.
4. Open the Network Control Panel: **System Preferences** > **Network**.



5. Create a new VPN connection:
   a. From the **Network** page, click the plus (**+**) icon in the left panel.

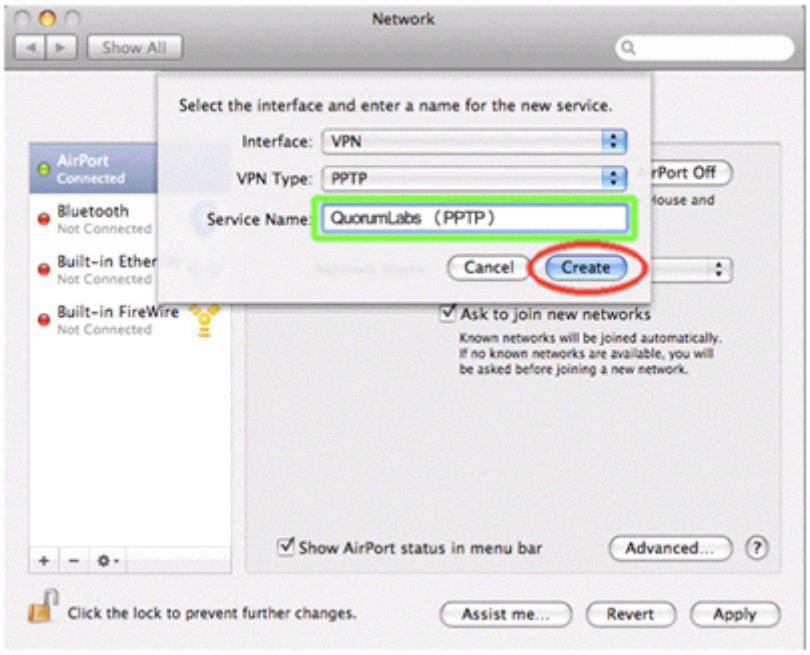
A new service sheet appears.

b. Click on the **Interface** popup menu to specify the properties for this new connection.

c. In the **Interface** drop-down list, choose `VPN`.

d. In the **VPN Type** drop-down list, choose `PPTP`.

e. In the **Service Name** field, name this new connection `Quorum (PPTP)`.



f. Click **Create**.

This new connection appears in the list of connections in the left pane.

6. Make a new Quorum configuration:

a. From the **Network** page, click the **Configuration** drop-down list, then **Add Configuration**.

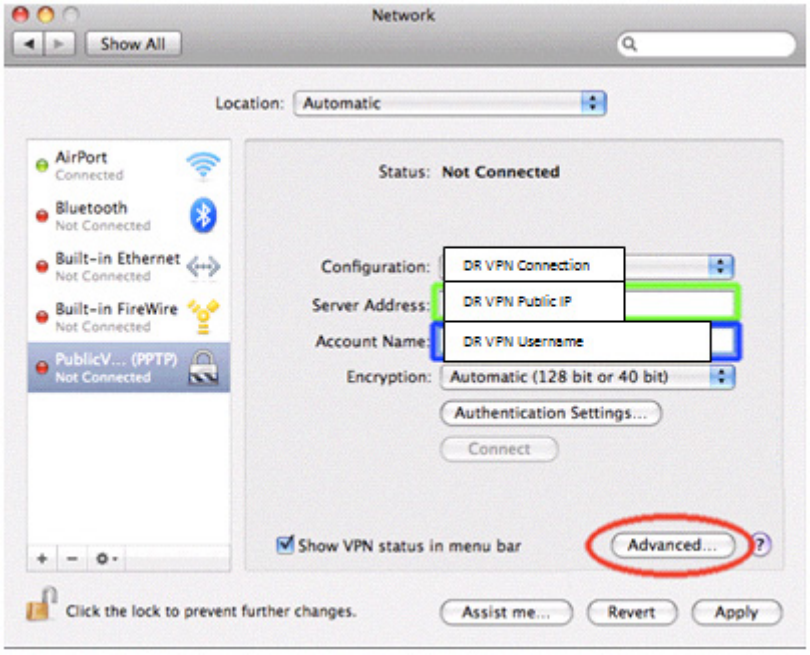
If you set up the VPN previously, the configuration appears in the list.

b. In the **Configuration** field, type `Quorum (PPTP)`, then **Create**.

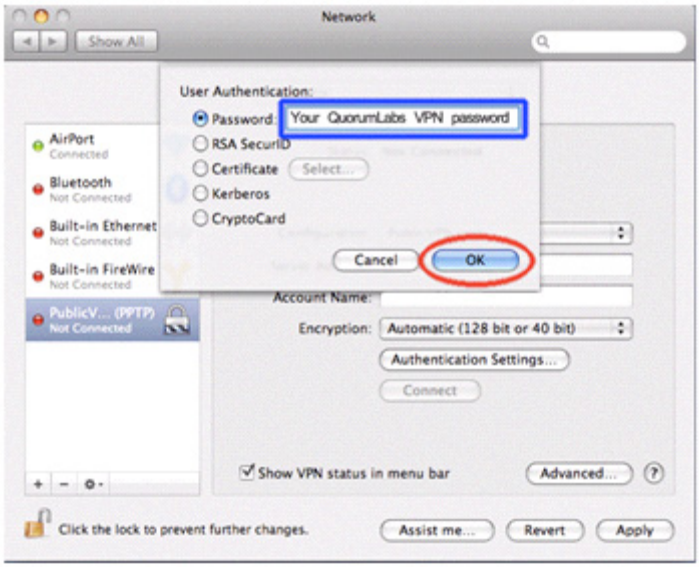
c. In the **Server Address** and **Account Name** fields, specify the Quorum IP address (WAN IP address) and Quorum username.

d. In the Encryption drop-down, choose **Automatic (128 bit or 40 bit)**.

e. Click on the **Authentication Settings...** button. The **User Authentication** dialog appears.



f. In the **Password** radio button field, specify your Quorum username, then **OK**.



7. Adjust the VPN settings:
   a. From the **Network** page, click the **Advanced...** button at the bottom of the page.
   b. Select the **Send all traffic over VPN connection** check box. You must do so in order for the connection to work effectively.
   c. (Recommended) Select the **Use verbose logging** check box. This option makes it easier to see additional information in the event that you have connection problems.
   d. Click **OK**.
8. Save the configuration. Click **Apply** in the bottom right of the Network page.
9. Establish the VPN connection.
10. Verify that the DR Appliance is up. From a command line prompt, ping the DR Appliance's LAN IP address.
11. Log on to the DR Appliance's web portal. From a web browser, go to `https://<DR-Appliance-LAN-IP-Address>`.

    You should now have full access to the RNs and File Level Backups on the DR Appliance.

Next Step: Step 4: Connect the Appliance to your network

3.2.4 Step 4: Connect the Appliance to your network

To connect the onQ Appliance:

1. Disconnect the standalone computer and connect the onQ Appliance to your target network using the active RJ-45 connection.
2. Verify that you can launch the onQ Portal.

   After the onQ Appliances reboot, you should be able to launch the onQ Portal from a browser anywhere on your network to the onQ local and onQ remote addresses (or names, if DNS is properly configured) that were configured above.
3. Reboot the onQ Appliance.

Next Step: Step 5: Change hypervisor's NTP server

3.2.5 Step 5: Change hypervisor's NTP server

Your onQ Appliance's hypervisor uses a default NTP server (`clock.fmt.he.net`). Set the NTP server to a local NTP server. To change the hypervisor's settings, go to Configure Appliance's hypervisor settings.

Next Step: Step 6: Enroll protected nodes

Related Topics

Synchronize system time

3.2.6 Step 6: Enroll protected nodes

Each protected node must be enrolled with the HA Appliance. For instructions, go to:

- (Agent-based Windows PNs) Enroll protected nodes
- (Agent-based Linux PNs) Enroll protected nodes
- (Agent-less Linux/Windows PNs) Enroll protected nodes
- Enrollment in Windows Cluster Services Environment

When you log on to the HA Appliance's onQ Portal, don't forget to use the new password that you set in Step 4 of Step 2: Configure the onQ Appliances.

Next Step: Step 7: Configure RN startup dependencies

3.2.7 Step 7: Configure RN startup dependencies

If any of your recovery nodes depend on other recovery nodes (for example, your mail server depends on a Domain Controller), configure startup dependencies for such nodes. For instructions, go to "Configure startup dependencies" on page 264.

Next Step: Step 8: Start automated protection process

3.2.8 Step 8: Start automated protection process

Now you're ready to start the automated protection process on both the HA Appliance and the DR Appliance. For instructions, go to Start node protection.

3.3 Initial Backup

When you first start protection on your nodes, onQ initiates a full disk image of the disks' entire contents, including operating system, configuration, applications, and user data. Performing initial deduplication and storing the data on the HA Appliance takes time.

Transferring gigabytes or terabytes of data over a relatively low-bandwidth WAN can be time consuming even with onQ's integrated deduplication and compression.

One way to minimize the time consumed by this process is to take advantage of onQ's exclude list feature, which enables you to specify folders and files that do not need to be backed up. See Edit backup exclude list for details.

After the initial cycle, or "base image", the onQ backup process is always incremental, meaning that subsequent backups involve only new files or only the parts of large files that have changed. Subsequent backups involve far less data and are generally quite fast.

If practical, Quorum suggests that you perform the initial backup with both your HA Appliance and DR Appliance on the same LAN. You can also use a removable drive for the initial synchronization as outlined in Initialize the DR Repository.

The subsequent backups begin based on the **Window Start** time, which you specified when you configured your PN's **Backup Schedule** (see "Schedule backups" on page 217).

You can monitor progress of the backups. See Monitor protected nodes and Monitor DR Appliance. If you receive any email alerts, follow the instructions in Backup Alerts.

You know that the initial backup has finished on an HA Appliance by checking the **Backup Status** field on a **PN Status** page, and on a DR Appliance by checking the **RN Status** field on the **DR Status** page. After the initial DR backup completes, you can shut it down and move it to its permanent locale.

Related Topics

Backup and Recovery Workflow

Initialize the DR Repository
Initiate immediate backups

3.4 Initialize the DR Repository

Incremental updates between the HA Appliance and DR Appliance are extremely efficient due to the deduplication and compression techniques employed by onQ. The initial "seeding" of the DR repository, whether on-premise or Hybrid Cloud, typically requires that the entire (deduplicated, compressed) images of all PNs be transferred from the HA Appliance to the DR Appliance. This can be done in one of the following ways:

- **LAN Synchronization:** This is the fastest way. Place the DR Appliance on the LAN with the HA Appliance until they are synchronized, then relocate the DR Appliance to the secondary site. You know the appliances are synchronized when the timestamps for the `Last Update Time` match. (see Monitor DR Appliance). This method is simple and fast, but it is not always practical due to geographical or other logistical considerations. DR Appliance might need to be reconfigured if its IP configuration needs to be changed for installation at the remote site.
- **WAN Synchronization**: While this is the simplest way, and allows you to initialize your installation with the onQ Appliances in their final locations, it might not be practical depending on the ratio of WAN speed to the size of the repository. The table below shows estimates of synchronization times for various data volumes and WAN bandwidth allocations (onQ allows you to throttle the amount of WAN bandwidth used). These calculations estimate 15% deduplication/exclusion savings on the base image and a 2.5 compression ratio.

ESTIMATED HOURS TO SYNCH HA-DR

		Allocated megabits / sec					
		1.5	**3**	**5**	**10**	**100**	**1000**
GB of PN DATA	**100**	56.7	28.3	17.0	8.5	0.9	0.1
	250	141.7	70.8	42.5	21.3	2.1	0.2
	500	283.3	141.7	85.0	42.5	4.3	0.4
	1000	566.7	283.3	170.0	85.0	8.5	0.9

- **Export/Import using External Drive:** Export the HA repository to portable storage and then import it to the DR repository, if you have too much data for a LAN or WAN synchronization.

Related Topics

Backup and Recovery Workflow
Export or import Repository
Resize protected node's vdisk

4

The Basics

- Log on to Appliance's user interface
- Establish a VPN connection to Hybrid Cloud
- Log off from Appliance's user interface
- Drop-down menu
- Reboot onQ Manager
- Reboot Appliance
- Shut down or Restart Appliance
- About onQ Proxy
- Modify Appliance's network settings
- Configure Appliance's hypervisor settings
- Unlock UI
- (Agent-based PNs) Restart the onQ Service
- Modify protected nodes
- Delete protected nodes
- Protected Node Parameters

4.1 Log on to Appliance's user interface

You need a supported browser (see Browser Support) to access the onQ Appliance's user interface.

- onQ is delivered with pre-defined user names and passwords for `Admin` and `VARAdmin`. If you did not change the default passwords as recommended in Step 2: Configure the onQ Appliances, use the defaults outlined in Change user passwords.
- Each onQ user is pre-assigned one of three roles: *Administrator*, *Operator*, *VARAdmin* or *Restore.* Each role implies a different level of permission to access certain UI features. For more information, go to Add users.

To log on to the onQ Appliance:

1. (Hybrid Cloud Only) Initiate a VPN connection to Hybrid Cloud.
2. Ensure that:
   - JavaScript is enabled on your browser.
   - the onQ Appliance is in your browser's Trusted Sites list.
3. Point the browser to the onQ Appliance's LAN IP address. You assigned this IP address during the configuration process. See Step 2: Configure the onQ Appliances.

Note: Alternatively, you can access a DR Appliance from the HA Appliance's page footer links:



4. Type in your user credentials.

4.2 Establish a VPN connection to Hybrid Cloud

How you establish a VPN connection to Hybrid Cloud depends on your operating system.

To establish a VPN connection to Hybrid Cloud (Windows 7):

1. From the **Network and Sharing Center**, right-click on **DR VPN Connection**, then **Connect**.
2. Type the credentials that your Quorum Support Engineer provided you, then **Connect**.
3. Wait a few seconds for your client to establish a VPN Connection to the Quorum Data Center.

To establish a VPN connection to Hybrid Cloud (Windows XP):

1. From **Network Connections**, right-click on **DR VPN Connection**, then **Connect**.
2. Type the credentials that your Quorum Support Engineer provided you, then **Connect**.
3. Wait a few seconds for your client to establish a VPN Connection to the Quorum Data Center.

   If authentication is successful, the Connection window minimizes to your taskbar as a dual computer icon. Run your mouse over it to see the status of the connection.

To establish a VPN connection to Hybrid Cloud (Mac OS X):

1. Do one of the following:
   - Choose the VPN configuration from the VPN menu.
   - From the **Network** page, select your VPN connection from the connection list, then click the **Connect** button.
2. Wait a few seconds for your client to establish a VPN Connection to the Quorum Data Center.

4.3 Log off from Appliance's user interface

Logging off closes the user interface. You will need to log on again to restore the UI.

To log off:

1. Select the drop-down menu.
2. Choose **Log Off**.

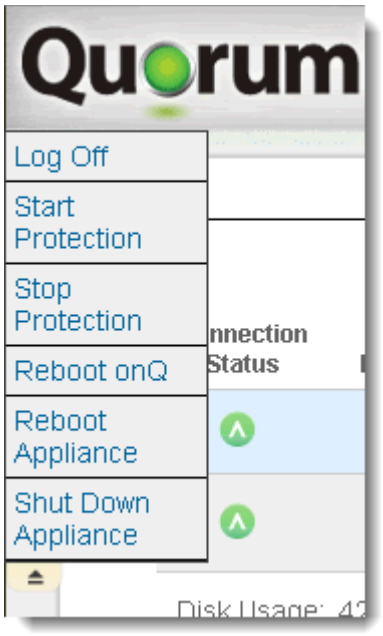
4.4 Drop-down menu

If you examine the left margin of the display, in the space to the left of the tabs and below the letter Q in the Quorum logo, you will discover a tiny icon that includes a representation of an arrowhead pointing downwards.



You can see this icon regardless of which major tab you have selected. When you click the icon, a pull-down menu appears as shown below:

The **Change Password** option does not display for users assigned to the `Administrator` role because such users can perform this task via the **USERS** page, unlike users assigned to the `Operator` or `Monitor` roles, which only have the **Change Password** and **Log Off** options.



When you open the menu, the drop-down icon migrates to the bottom of the menu and the arrowhead points upward. Click the icon again to close the menu.

The menu offers the following five critical actions. These are actions that must be selected cautiously. They can have undesired consequences if invoked accidentally. The tiny drop-down icon helps avoid unintended choices.

Log Off

Start Protection

Stop Protection

Reboot onQ Manager

Reboot onQ Appliance

Shut Down onQ Appliance

4.5 Reboot onQ Manager

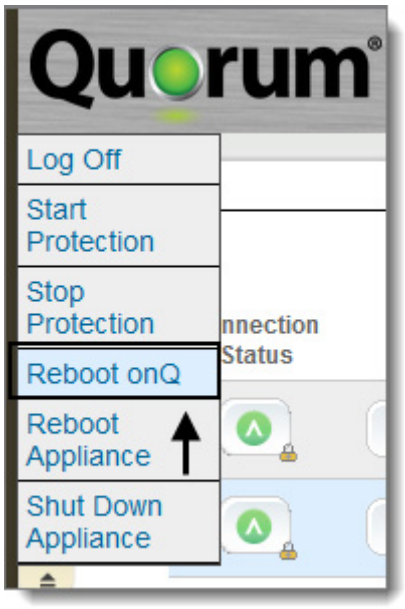
Some software upgrades require that you reboot the virtual machine on which the onQ Manager depends. In doing so, all processes managed by the onQ Manager are restarted. The underlying hypervisor, however, is not affected and neither are any RNs running in production or test mode.

A reboot temporarily stops protection. However, onQ Manager returns the onQ Appliance to its original protection state after the reboot:

- If protection was *on* prior to the reboot, onQ Manager turns on protection after the reboot.
- If protection was *off* prior to the reboot, onQ Manager does not turn on protection after the reboot.

To reboot the onQ virtual machine:

1. Log on to the onQ Appliance's onQ Portal.
2. Select the drop-down menu.
3. Choose **Reboot onQ**.



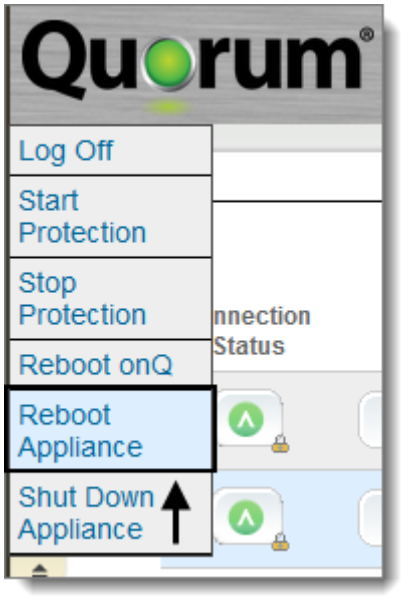
4. Click **Continue** to initiate the reboot.
5. After the onQ Manager reboots, verify that protection is on.

4.6 Reboot Appliance

Use this procedure to restart the entire onQ Appliance. The hypervisor and all processes managed by the hypervisor are restarted. Rebooting the onQ Appliance interrupts all protection operations and stops all recovery nodes.

To reboot the onQ Appliance:

1. Log on to the onQ Appliance's onQ Portal.
2. Select the drop-down menu.
3. Choose **Reboot Appliance**.



4. Click **Continue** to initiate the reboot.

   After the onQ Manager reboots, don't forget to restart protection.

4.7 Shut down or Restart Appliance

Use this procedure to immediately halt the onQ Appliance. The hypervisor and all processes managed by the hypervisor are halted. Shutting down the onQ Appliance interrupts all protection operations and stops all recovery

nodes.

To shut down the onQ Appliance:

1. Log on to the onQ Appliance's onQ Portal.
2. Select the drop-down menu.
3. Choose **Shut down** onQ Appliance.

To restart the onQ Appliance:

1. Ensure that the power is off, then turn it on (power-cycle).
2. Restart protection.
3. Restart any recovery nodes that were in use before the reboot occurred, if they don't start automatically.

   Virtual machines hosted on a hypervisor (specifically XenServer) start automatically when the onQ Appliance reboots. However, sometimes a virtual machines do not start up automatically upon reboot of the onQ Appliance.

4.8 About onQ Proxy

PNs need the onQ Proxy, not to be confused with onQ Central Proxy and PN Proxy or Proxy Host, to communicate with onQ for backups. PNs can share an onQ Proxy (aka *global* onQ Proxy) or each PN can have a separate onQ Proxy.

If you do not specify an onQ Proxy for the PN, onQ uses the global onQ Proxy. If you do not specify a global onQ Proxy, onQ uses its own IP address.

If you have any PNs on a different subnet, such as in a DMZ, you must specify an onQ Proxy on the PN itself. However, at the time of enrollment, the PN does not have a configuration so you cannot set its onQ Proxy until after enrollment. Therefore, before enrollment of a PN on a different subnet, you must temporarily set the global onQ Proxy to the onQ Proxy for that PN. After enrollment, revert to the global onQ Proxy.

4.9 Modify Appliance's network settings

You can change a variety of settings related to your onQ Appliance, including network settings:

- Fully Qualified Host name
- onQ Role (see Change Appliance's role)
- IP Address
- Subnet Mask
- Default Gateway
- onQ Proxy Address (see About onQ Proxy)
- Preferred DNS Server
- Alternate DNS Servers
- onQ Appliance's role
- Time zone

To modify an onQ Appliance's network settings:

Renaming of onQ is rare, but there are circumstances that might require that you to do so:

- Your company name or location changed, and your host name needs to realign with this new identity.
- You initially deployed one onQ, and now you're adding another.
- You are replacing your onQ Appliance so as to upgrade the hardware. There cannot be two onQs with the same name. In this case, you would have two onQs with the same name until after the data from old onQ Appliance migrates to new onQ Appliance.

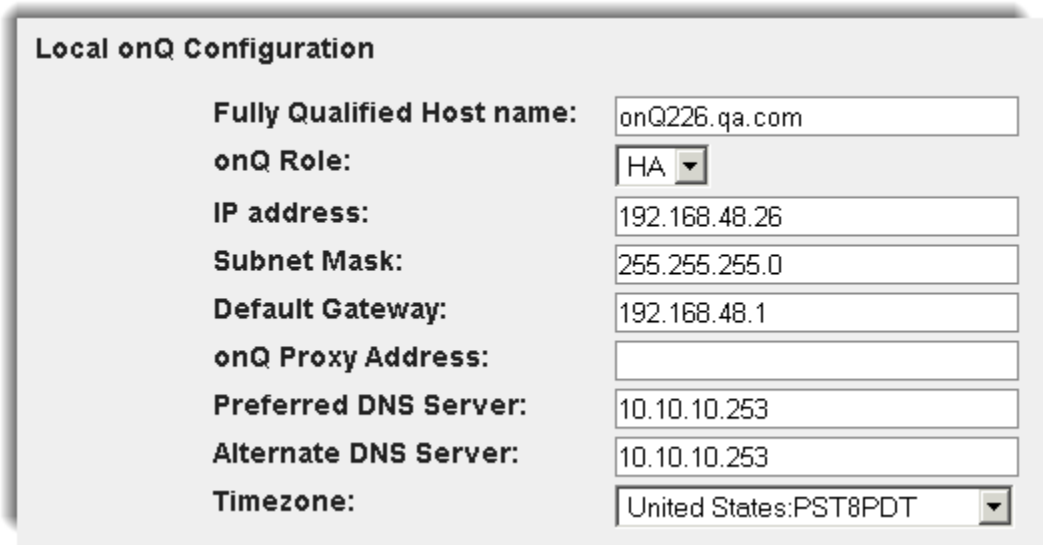
In the case of an agent-less PN enrollment, if you need to rename your onQ, do one of the following:

- If you have only one PN: (1) From the onQ Portal, delete the PN (Delete protected nodes), choosing to retain ...the associated data from the repository and corresponding RN.... (2) Rename the onQ. (3) Re-enroll the PN as outlined in (Agent-less Linux/Windows PNs) Enroll protected nodes.

- If you have multiple PNs: (1) Rename the onQ; (2) On each ESXi/ESX server, delete from disk all PN proxies that use the old onQ host name. (3) Re-enroll the PNs as outlined in (Agent-less Linux/Windows PNs) Enroll protected nodes.

1. Log on to the onQ Appliance's onQ Portal.
2. Go to **APPLIANCE CONFIG** tab > **onQ (LOCAL)** page > **MODIFY** button.

   The **Modify Local onQ Setup** page appears:



3. Make your change, then click **SAVE**.

   Click **REVERT** to return the fields to their initial values, or click the cancel icon (at the upper left) to discard your changes and return to the previous page.

4.10 Configure Appliance's hypervisor settings

You can configure a variety of settings related to this onQ Appliance's hypervisor, including:

- Hypervisor name
- NIC Speed (not configurable: onQ Portal reports the speed that's detected)
- Hypervisor IP Address

- Hypervisor Password
- Subnet Mask
- Default Gateway
- Preferred DNS Server
- Alternate DNS Server
- NTP Server (default is `0.us.pool.ntp.org`)
- Time zone
- iDRAC (see Change iDRAC network configuration)

To configure an onQ Appliance's hypervisor settings:

Your onQ Appliance's hypervisor is XenServer. The onQ Appliance ships with the hypervisor information pre-configured.

The hypervisor name is based on a Quorum naming convention: Hyper-Q-*CustomerName-ApplianceType*. This naming convention provides you the best customer support possible because it communicates the onQ Appliance's role and your topology.

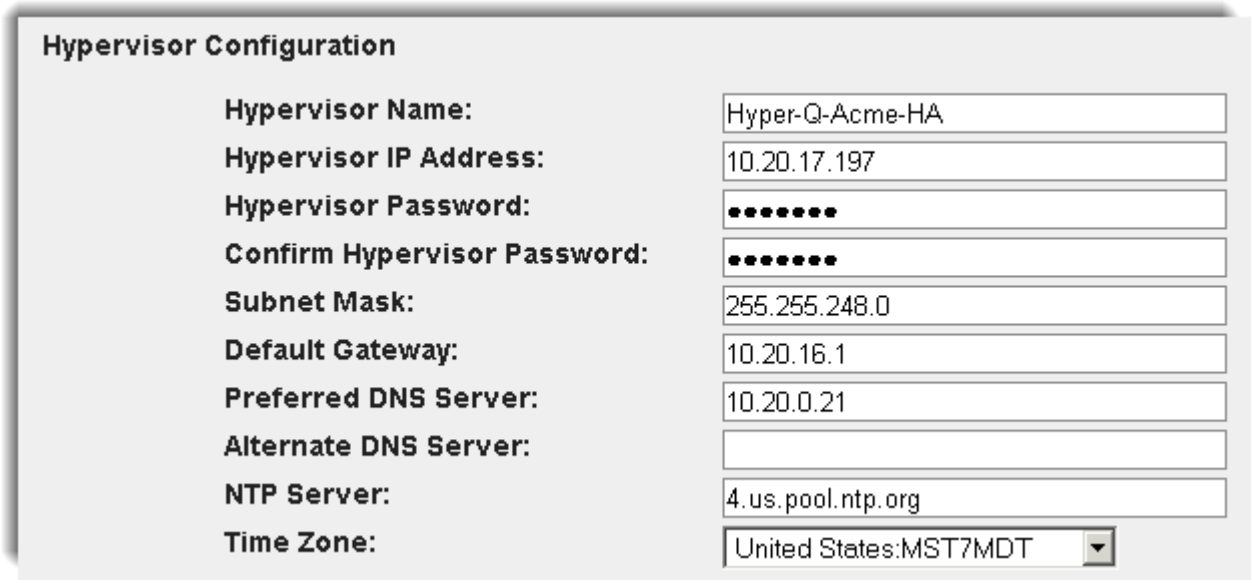
Quorum recommends that you retain the pre-configured hypervisor name; however, Quorum recognizes that your company might have its own naming conventions. Before you change the hypervisor's name, contact Quorum Support with your naming convention requirements so that we can rename the hypervisor. without losing sight of the relationships between onQ Appliances.

Hyper-Q-Acme-HA	Acme's HA Appliance.
Hyper-Q-Acme-DR	Acme's DR Appliance.
Hyper-Q-Acme-DRaaS	Acme's Hybrid Cloud.
AV-Acme-DR	Acme's AV Appliance, which is an archive of the DR Appliance.
AV-Acme-HA	Acme's AV Appliance, which is an archive of the HA Appliance.
Hyper-Q-Acme-MT1	One of Acme's multi-tenant onQ Appliances. Multi-tenancy is identified by the Configuration field under **Dashboard** tab > **onQ Status** page.

1. Log on to the onQ Appliance's onQ Portal. You must log on as `varadmin` user.

2. Click the **APPLIANCE CONFIG** tab > **HYPERVISOR** page > **MODIFY** button.

   The **Modify Hypervisor Config** page appears.



3. Make your changes, then click **SAVE**.

   Click **REVERT** to return the fields to their initial values, or click the cancel icon (at the upper left) to discard your changes and return to the previous page.

4. Reboot the onQ Appliance.

4.11 Unlock UI

The onQ Portal has pages that have icons that lock and unlock. This feature ensures that you don't make an unintentional change to your protection configuration. Only users of role *Administrator* or *Operator* can unlock these icons; users of role *Monitor* cannot unlock these icons.

After you unlock the icons they become command buttons, enabling you to perform certain actions. To unlock these icons, simply use the padlock icon

to toggle between these two states:

	click to enable UI command buttons
	click to disable UI command buttons

For example, if the padlock icon shows a closed padlock and is labeled **Unlock to enable actions**, you can view the status information related to protected nodes, but you cannot perform any actions.

To unlock the UI for a given page:

1. Go to the page you want to unlock:
   - (HA) **DASHBOARD** tab > **PROTECTED NODES** page
   - (DR) **DASHBOARD** tab > **DR STATUS** page
   - (HA and DR) **DASHBOARD** tab > **RECOVERY NODES** page
2. Locate the closed padlock icon in the lower left corner.
3. Click on the icon to switch the icon to the open state, assuming the icon is in the closed state.

   The icon changes to an open padlock and is unlabeled. Some of the fields in the table are revealed to be clickable buttons:

4.12 (Agent-based PNs) Restart the onQ Service

If the onQ Service is not running on the protected node, the HA Appliance cannot establish a secure connection to that protected node. Restarting the onQ Service during a backup or upgrade can cause backup/upgrade failures.

(Windows) To restart the service:

1. RDP to the PN.
2. Go to:

   Start > **Programs** > **Control Panel** > **Services**

   OR

   Start > **Administrative Tools** > **Services**
3. Right-click on **Quorum onQ Monitor** (aka onQ Service), and choose **Restart**.

(Linux) To restart the service:

1. Verify that the service is installed:

```
> rpm -qa | grep node
```

2. Start the service:

```
> service dcrm-node start
```

4.13 Modify protected nodes

During the setup process you identified the nodes you wanted to protect. From time to time you might need to change the settings on these nodes.

To edit a protected node (HA Appliance):

1. Log on to the HA Appliance's onQ Portal.
2. Click the **PROTECTION CONFIG** tab.
3. Select the node that you want to modify, then click **MODIFY**.

   The Modify a Protected Node page appears.
4. Make your changes to the node parameters. You cannot change the **Hostname** field from the **MODIFY** window.
5. Click **SAVE** to save your changes and return to the **PROTECTION CONFIG** tab.

   A newly modified protected node appears in the list.
6. Restart protection, if applicable.

To edit a protected node (DR Appliance):

1. Log on to the DR Appliance's onQ Portal.
2. Click the **PROTECTION CONFIG** tab.
3. Select the node that you want to modify, then click **MODIFY**.

   The Modify a Protected Node page appears.
4. Specify values for the node parameters.

   You cannot change the **Hostname** field from the **MODIFY** window.
5. Click **SAVE** to save your changes and return to the **PROTECTION CONFIG** page.

   A newly modified protected node appears in the list.

4.14 Delete protected nodes

During the setup process you identified the Protected Nodes (PN) that you want to protect.

You might want to delete a specific PN if:

- You removed the host from production. In this case, you want to "free up" a license or vdisk disk space, especially if you don't have enough disk space to enroll protected nodes in production. Licensing is based on the number of nodes being protected.
- You want to move the PN to another HA Appliance. Only one HA Appliance can protect a given PN.

The onQ Portal will not let you delete a PN if:

- PN is in production.
- PN is being backed up.
- HA Appliance's RN is in test mode, production mode, or in the process of being tested (self-test) or updated.
- DR Appliance's RN is in the process of being tested (self-test) or updated.

However, if you think you might have plans for the PN in the future, disable it instead.

To delete a PN:

After you delete the PN, you cannot start the RN on the HA Appliance. Also, When you delete a PN on an HA Appliance, onQ automatically propagates that deletion on the DR Appliance.

1. Log on to the HA Appliance's onQ Portal.
2. Ensure that the HA-to-DR link is working. Go to Test HA-to-DR Link.
3. Disable protection for the PN that you want to delete, then do the following if applicable:
   - If PN backup is in progress, either wait for it to complete or stop the backup.
   - If RN update is in progress, either wait for it to complete or stop the update.
   - If RN is running a self-test, wait for it to complete.

- If RN is running in test or production mode, power off the RN.

4. Click on the **PROTECTION CONFIG** tab, then select the PN from the list.
5. Click the button labeled with the minus sign (**-**).

Note: Deleting a PN purges the PN's virtual machine (vdisk); a deletion does not purge all the data associated with that PN.

6. (Optional) Select the **Remove**... check box to delete all the PN's data, then **Yes**. Alternatively, delete the orphan data later.

   The onQ Portal returns you to the **PROTECTION CONFIG** tab, and the PN disappears from the list.

4.15 Protected Node Parameters

Specify these parameters when you want to change a protected node's configuration from the **PROTECTION CONFIG** tab.

Basic Parameters

- **Hostname:** If you are adding a protected node, type the host name or IP address of the node to protect. This field is display-only under **MODIFY**.
- **onQ Proxy Service Rev**: (This parameter is not applicable for agent-based enrollment.) The software revision running on the PN proxies.
- **Proxy Host**. (This parameter is not applicable for agent-based enrollment.) This is the ESX/ESXi server that's hosting the PN proxies.
- **onQ Node Service Rev**: (This parameter is not applicable for agent-less enrollment.) The software revision running on the PNs.
- **Backup Mode**: (This parameter is not applicable for agent-based enrollment.) As outlined in "(Workflow) Fail over a PN to an RN" on page 272 and "(Workflow) Fail back an RN to a PN" on page 274, this parameter reflects whether onQ is currently backing up the agent-less PN (**`Agentless`**) or the RN (**`Recovery Node`**) in production mode. onQ Portal changes this parameter automatically as it powers on and off the RN.
- **Group Name**. A logical means of grouping your PNs. Groups are required to configure startup dependencies as outlined in Configure startup dependencies. For Agent-less PNs, the onQ Portal groups the PNs by ESXi server. If you do not want these PNs in such a group, clear the default Group Name field at the time of enrollment to remove the PNs from this group and place them in the shared pool.
- **Backup Schedule**. onQ performs an initial backup of your PNs when you first start protection. Modify the default backup schedule to instruct onQ when and how often to perform subsequent backups. Go to Schedule backups.
- **Backup Retention Time**. The maximum number of days to retain backed-up data that is no longer referenced. Go to Change Backup Retention Policy.
- **File Backup Only**: When set to **Yes**, squirtcopy backs up up only those files and directories that you specify in the **Backup Include List** (see "Advanced Parameters" on page 60). This parameter can only be enabled

by Quorum Support and is typically used only for file servers. Note the following when this parameter is enabled:

  - **Auto RN Creation?** parameter is hidden because it is not applicable.
  - **Recovery Node Configuration** pane is hidden because the protected drives are determined via the Include List.

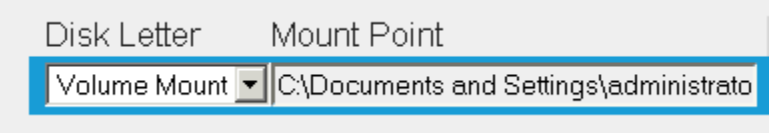
- **OS Type?**: Both Linux and Windows operating systems are supported, though the parameters available vary.
- **Auto RN Creation?**: This policy instructs onQ to create a recovery node for this protected node, and when it should be available. This policy represents your potential downtime in the event of a disaster (aka RTO). See also "Modify recovery node creation policy" on page 184.
- **Disable Protection?**: Select **No** to enable or **Yes** to disable protection for this node. If you disable protection, the most recent re-build of the recovery node is retained but backups are suspended. If you disable protection while a backup is in progress or an RN is being rebuilt, that process continues to completion but further backups are suspended. When you re-enable protection, backups resume and subsequent RN rebuilds are based on the existing RN. See "Disable protection of nodes" on page 189.
- **Disable Replication to DR?**: Select **No** to enable or **Yes** to disable replication for this node. You might want to disable replication if this node is not important for business continuity, but you do want backups and the ability to archive them. If you are unsure, keep this setting enabled. See Disable replication for individual nodes.

Recovery Node Configuration Parameters

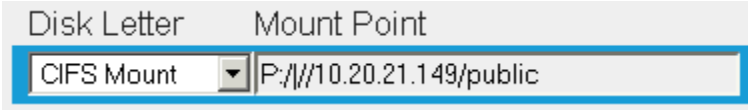
There are a few recovery node parameters that vary slightly depending on the operating system:

- **Number of CPUs**: # of CPUs for this recovery node.
- **RN Memory**: RN's memory in GBs.
- **Disk Letter**: Disk's letter, Volume Mount, or CIFs Mount.
- **Mount Point**: This parameter varies by platform:
  - (Linux) Where the protected node's filesystem is mounted.
  - (Windows) Full path to the directory in which this drive is mounted. Windows drives can be mounted to appear as a directory in another disk. These mounted drives must be protected separately:

1. Select **Disk Letter** drop-down box > **Volume Mount**.
2. In the Mount Point text box, provide the path to the directory. For example: `C:\Documents and Settings\administrator\My Documents\public`.



- (Windows) Full path (***<mount-point>|<url>*** where `url` is ***//<ip-address>/<path>***) to the mounted shared volume:
  1. Select **Disk Letter** drop-down box > **CIFS Mount**. For example: `P:/|//10.20.21.149/public`.



  2. In the Mount Point text box, provide the path to the directory.
  3. Set the necessary privileges as outlined in Enable onQ to back up shared volumes.

- **Format**: File system type. This parameter is for Linux only.
- **VDs**: # of virtual disks required to hold the configured disk. Every multiple, or portion thereof, of 1980GB requires a virtual disk. For Unix RNs, each configured disk must fit on a single VD: spanning of virtual disks is only supported for Windows RNs.
- **Size**: Size of the disk in GBs.
- **+**: Adds a new line to the Recovery Node Configuration list.
- **-**: Removes a highlighted line from the Recovery Node Configuration list.
- **REVERT** button 

: Reverts all the fields to the last saved values.

Advanced 

Parameters

There are several **ADVANCED** parameters:

- **Execute Pre Snapshot script?** (This parameter is not applicable for agent-less enrollment.) Executes before a snapshot starts on the protected node. See "Run custom backup scripts" on page 237.

- **Execute Post Snapshot script?** (This parameter is not applicable for agent-less enrollment.) Executes soon after a snapshot completes on the protected node. See “Run custom backup scripts” on page 237.
- **Enable Filter Driver?** (This parameter is not applicable for Linux-based protected nodes and agent-less enrollment.) See “(Agent-based PNs) Methods of Performing Incremental Backups” on page 241.
- **Truncate SQL logs after backup?** (This parameter is not applicable for Linux-based protected nodes.) After a successful backup, onQ truncates the SQL logs.
- **VSS Copy Only?** (This parameter is not applicable for agent-less enrollment.) onQ’s backup utility clears VSS flags after it performs backups. This process can conflict with some applications (for example, Backup Exec) that depend on these VSS flags. For example, other backup products can trigger Exchange and SQL Server to truncate log files after incremental backups. Only change this setting to **Yes** if you are using an application that relies on and updates these VSS flags; otherwise, set this parameter to **No** (default). Setting this parameter to **Yes** ensures that onQ does not alter or reset VSS flags and leaves your files, generally transaction logs, as-is.
- **PN Scan Threads:** This is the number of threads used during the scan process. The more threads used, the more likely you are to lock up your PN. If you are unsure, specify 1 thread.
- **PN Transfer Threads**: The number of threads used during the backup process. The more threads used, the sooner the backup process completes. Ensure that the client machine has enough resources to handle the number of threads you specify. If you are unsure, specify 4 threads.
- **CPU Resource Limit:** Sets the CPU usage limit for the backup process on the protected node during the transfer phase. The default is 30% of all cores.
- **Disk Resource Limit:** Sets the disk bandwidth limit for the backup process on the protected node during the transfer phase. The default is 170 MB/second.
- **Network Resource Limit:** Sets the network bandwidth limit for the backup process on the protected node during the transfer phase. The default is 170 MB/second.
- **Enable Cluster Support?**: (This parameter is not applicable for Linux-based protected nodes. Also, for information on how to work with PNs/RNs in a cluster environment, refer to Enrollment in Windows Cluster Services

Environment.) If you select **Yes**, specifies that the node is part of an MSFT cluster (Microsoft Cluster Server/MSCS). In a nutshell, cluster support works by creating a virtual SAN to replicate your physical SAN. The PN accesses the SAN using a virtual IP address. For each volume, the onQ Appliance creates a virtual disk.

  - **Cluster Volumes**: PN's volumes/partitions. List of non-boot drives. From this list, choose the volumes that correspond to your cluster.
  - **Virtual SAN IP Address**: Any dedicated and unique IP other than the SAN's physical IP address. This IP address must be on the same sub-net as the onQ Appliance's physical IP address.
  - **Virtual SAN Gateway Address**: Virtual SAN's Gateway address.
  - **Virtual SAN Subnet Mask**: Virtual SAN's subnet.

- **onQ Proxy Address**: Specifies the onQ Proxy that the PN needs to use to communicate with the onQ for backups. Unless your PN is on a separate subnet, leave this field blank. For more information, go to About onQ Proxy.
- **Backup Exclude List**: The button shows the exclude list to be either **DEFAULT** (as defined by Quorum) or **CUSTOM** (as defined by your administrator). Click the button to modify the backup exclude list for this protected node. See "Edit backup exclude list" on page 210.
- **Backup Include List**: onQ backups up only those files and directories that you specify. This button appears if **File Backup Only** is enabled. See "Edit backup include list" on page 213.
- **Startup Dependencies**. onQ will start a given RN after the RNs on which it depends are up and running. See Configure startup dependencies.
- **RN Services**. Some services installed on a PN will cause problems if they attempt to run on the RN. In this case, you can configure the startup settings for these services. See Edit RN services list.
- **RN Networks**. onQ enables your RNs to run on multiple networks. See Create Custom Networks for RNs.

Related Topics

Modify protected nodes

(Agent-based PNs) Add protected nodes manually

(Agent-based PNs) Restart the onQ Service

5

Enrollment in non-cluster Environment

- (Agent-based Windows PNs) Enroll protected nodes
- (Agent-based Linux PNs) Enroll protected nodes
- (Agent-less Linux/Windows PNs) Enroll protected nodes
- (Agent-based PNs) Add protected nodes manually

5.1 (Agent-based Windows PNs) Enroll protected nodes

onQ supports both agent-based and agent-less PNs. Nodes that are not hosted by VMware require that an agent (onQ Service) be installed. Agent-based enrollment enables you to use all of the operation and monitoring features that onQ has to offer.

Each agent-based PN must be "enrolled" with the onQ HA Appliance using the **Protect Me** button and the onQ EZ Installer, which analyzes your PN and pre-configures it on the onQ Appliance.

Even if you added the new node manually and it appears in the list of protected nodes, you must enroll the node before the onQ Manager can protect it. Previously protected nodes that are still running the onQ Service do not need to be enrolled; in this case, simply add the nodes manually, then start protection.

This *Protect Me-based* enrollment instructs the PN to download and extract the `QuorumWinBCVSetup<bitVersion>-BCV<build>.ez` package from the onQ Portal's **Downloads** tab. This package is only used for Protect Me-based enrollments, not *manual* installs.

Warning: In almost all instances, Non-NTFS partitions are *not* supported for VSS-based snapshots. Do not protect such partitions with onQ. Protecting such non-NTFS partitions might result in the Quorum agent terminating the backup. The only exception is a volume configured with REFS, which is available in Windows Server 2012.

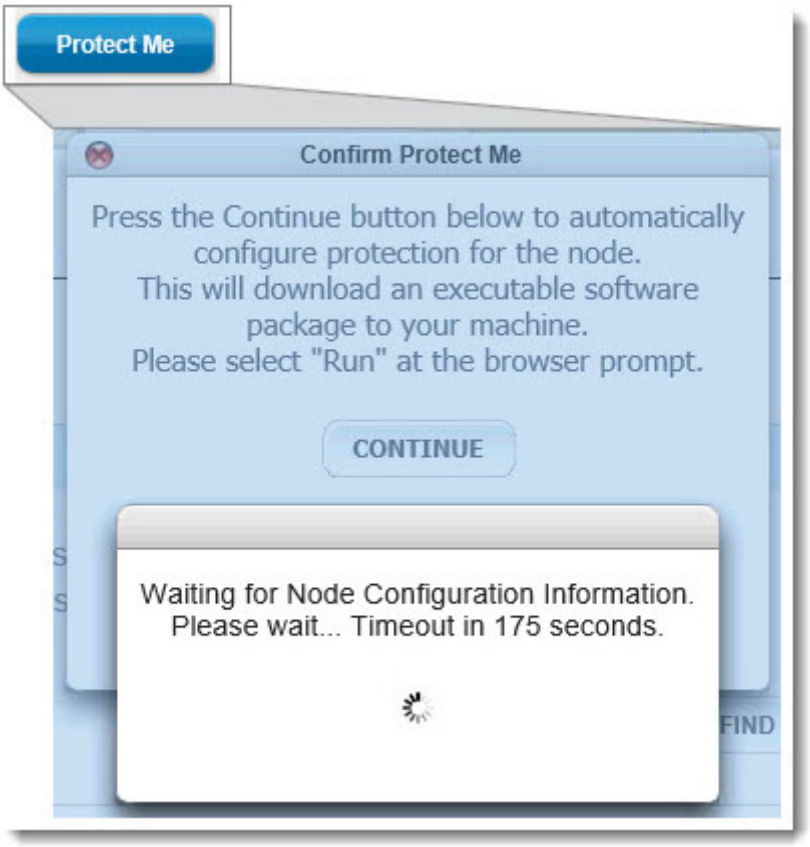
(Recommended) To enroll an agent-based Windows PN:

Use this procedure if you're trying to install an agent-based Windows PN for the first time. You must perform this procedure from the node that you want to enroll. Repeat this procedure for each node that you want to enroll.

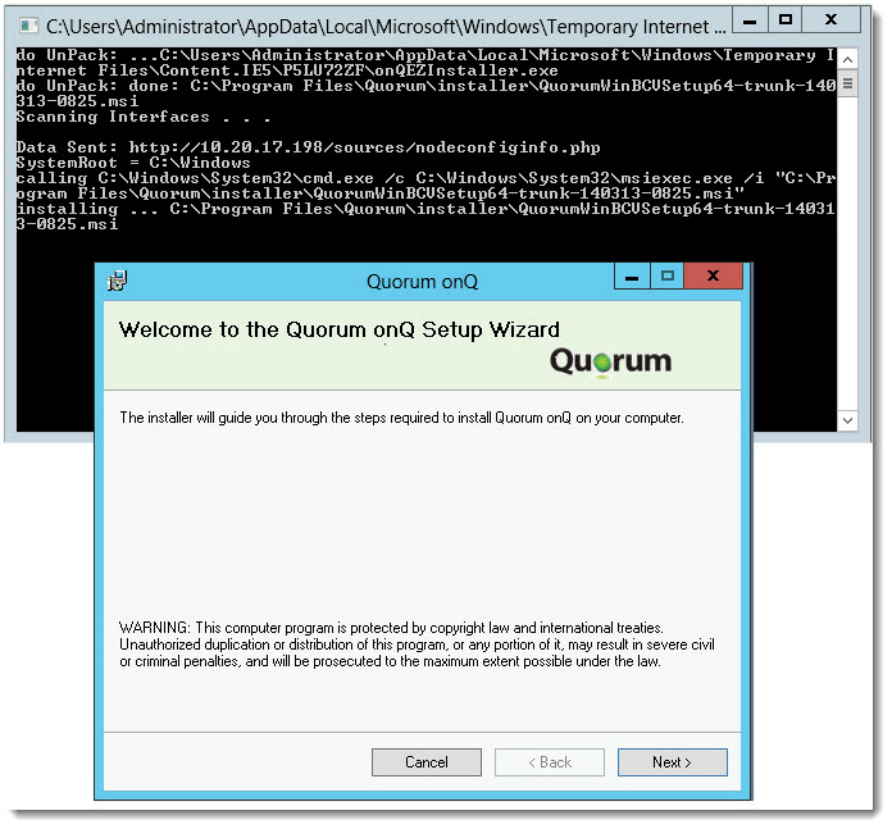
1. RDP to the server that you want to enroll as a PN, then log on as a user with administrative privileges.
2. From that server, launch a browser (see "Browser Support" in onQ Release Notes).

3. Log on to the HA Appliance's onQ Portal as `varadmin`.
4. In the onQ Portal, go to **PROTECTION CONFIG** tab > **Protect Me** button, then **CONTINUE**.

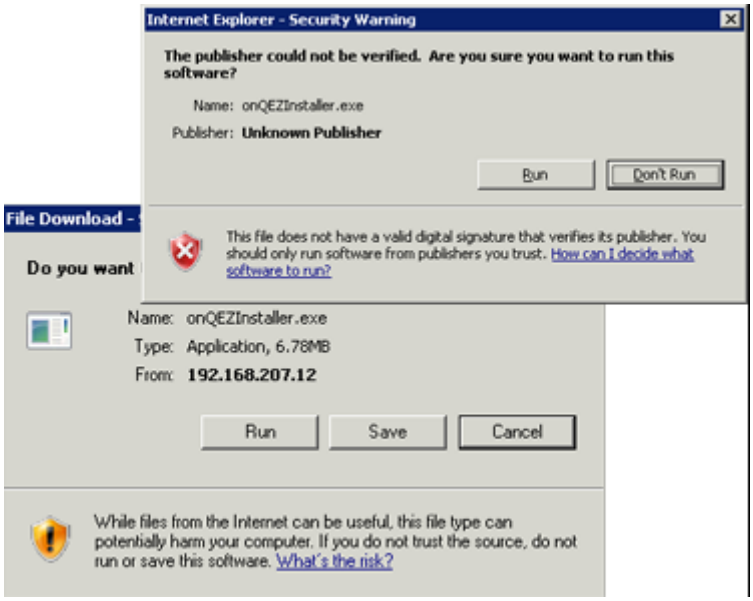
Note: The protection process instructs the PN to download and extract the `QuorumWinBCVSetup<bitVersion>-BCV<build>.ez` package from the onQ Portal's **Downloads** tab. This package is only used for *Protect Me-based* enrollments, not *manual* installs. The `.msi` installer only upgrades the onQ service component, while the `.ez` installer also enables enrollment.



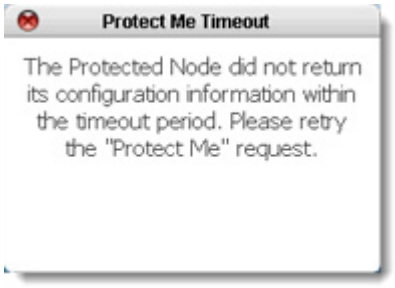
5. From your browser, run the onQ Service installer (`onQEZInstaller.exe`) to launch the **Setup Wizard**.



When prompted by your browser, click **Run**, the **Run** again to confirm the `Unknown Publisher`. If you're using Firefox, you need to double-click on the executable before you can run the installer.

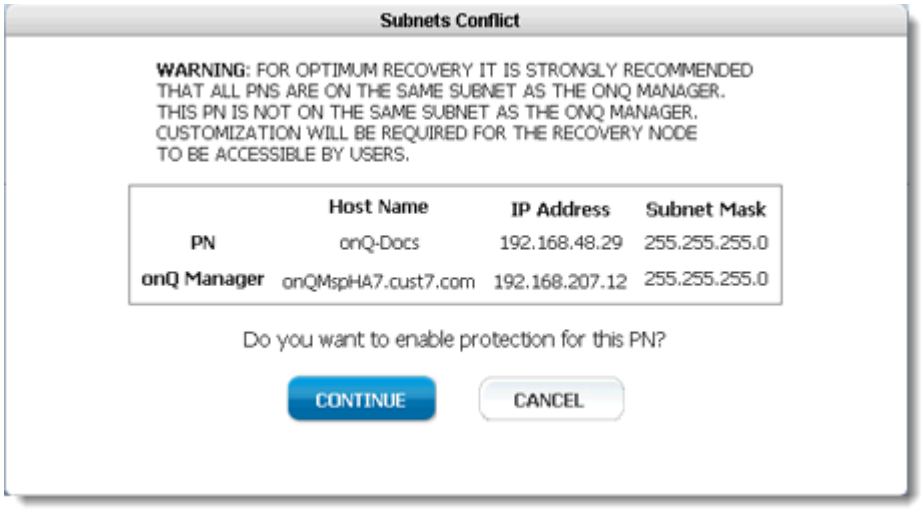


- If you don't install the service within 180 seconds, onQ returns the following error. In this case, simply close the window and repeat this procedure again.



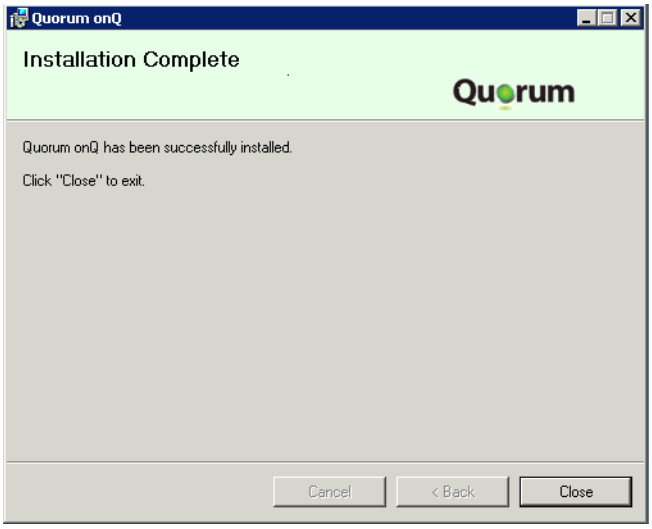
- If the installer detects that the onQ Manager and the PN are on

different subnets, the installer alerts you. Click **CONTINUE**.



6. Install the onQ Service by following the on-screen instructions in the **Setup Wizard**.

   Congratulations, you've installed the onQ Service.



7. Modify PN parameters.

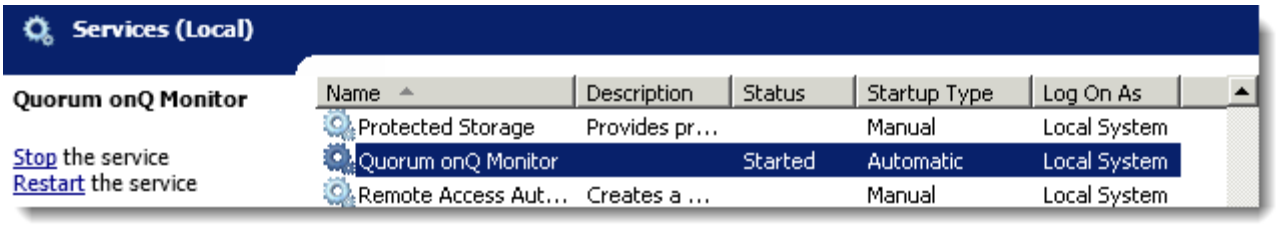
   The **Modify Protection Parameters** dialog presents you with a summary of the PN's proposed configuration.

   Specify values for the node parameters, then click **SAVE** to add the node to the list of protected nodes.

Optionally, you can click **SAVE** without modifications. You can modify the node at any time.

If the installer fails at this point because there isn't enough disk space, refer to (Agent-based PNs) Add protected nodes manually; otherwise, the PN appears in the list.

8. On the protected node, open the following ports in order for onQ to see the protected node.
   - UDP port `5990`
   - TCP ports `5000` and `5990`
9. Verify that the PN appears in the list of protected nodes.
10. Verify that the onQ Service is running on the PN:



(Alternative) To *manually* install or reinstall agent-based Windows PN software:

Use this procedure to:

- Manually install an agent-based Windows PN for the first time. Normally, there's no need to perform this procedure in this case because the Protect Me-based enrollment, as outlined in (Recommended) To enroll an agent-based Windows PN:, simplifies the installation process for you by retrieving the appropriate package from the onQ Portal's **Downloads** tab. However, you might need to perform a manual enrollment (install) if the initial Protect Me-based enrollment of the PN failed. Sometimes this happens. For example, old versions of Firefox cannot automatically select the correct operating system version (32-bit or 64-bit) for your PN.
- Reinstall (aka manually upgrade) node software on a PN that is already on onQ's protection list so as to fix a PN upgrade or connectivity failure. As outlined in (Agent-based Windows PNs) Update node software, the automatic upgrade process keeps your PNs up to date. However,

sometimes upgrade failures occur, requiring that you reinstall the node software. For example, a PN can go offline and miss an auto-update.

Note: There are cases where a **Protect Me** > **Setup Wizard** > **Repair** | **Remove** methods can resolve most PN upgrade problems, but uninstalling from the Windows Control panel and manually installing as outlined below corrects all issues.

1. **(Existing PNs Only)** If this is a previously enrolled PN, uninstall the existing node software on the PN: From the Windows Control panel, select the Quorum onQ Monitor service, then **Uninstall**.
2. RDP to the server that you want to protect, then log on to that server as a user with administrative privileges.
3. From that server, launch a browser (see Browser Support), then log on to the HA Appliance's onQ Portal as `varadmin`.
4. Download the onQ Service installer.
   a. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **DOWNLOADS** page.

Warning: Do not download the `.ez` file, which is only intended for Protect Me-based enrollments as outlined in (Agent-based Windows PNs) Enroll protected nodes, not manual installs. The `.msi` installer only upgrades the onQ service component, while the `.ez` installer also enables enrollment.

b. Scroll down and select the `QuorumWinBCVSetup<bitVersion>-BCV<build>.msi` file that matches your server's operating system.

Filename	File ID	Size	Timestamp	Type
QuorumWinBCVSetup64-BCV-3.8-140623-0716.ez	1232	7721011	2014-07-01 11:41:21	Node SW
QuorumWinBCVSetup32-BCV-3.8-140623-0716.msi	1233	6085120	2014-07-01 11:41:21	Node SW
QuorumWinBCVSetup64-BCV-3.8-140623-0716.msi	1234	6250496	2014-07-01 11:41:21	Node SW
DCRMnode-4.5.2BCV-6032.el5.i386.rpm	1235	3681271	2014-07-01 11:41:21	Node SW
QuorumWinBCVSetup32-BCV-3.8-140623-0716.ez	1231	7555635	2014-07-01 11:41:20	Node SW

c. Click the **DOWNLOAD** button, saving the file to your `Downloads` folder.

d. Open the `Downloads` folder. Execute the newly downloaded msi installer.

5. Follow the wizard's instructions to install the onQ Service on the PN.

Related Topics

(Agent-based PNs) Add protected nodes manually

5.2 (Agent-based Linux PNs) Enroll protected nodes

onQ supports both agent-based and agent-less PNs. Nodes that are not hosted by VMware require that an agent (onQ Service) be installed. Agent-based enrollment enables you to use all of the operation and monitoring features that onQ has to offer.

Linux PN preparation is very important for RN builds to succeed. In some cases, different RHEL release distributions require unique tasks:

- **install xen-aware kernel package**. Unlike with RHEL 5.x, RHEL 7 and RHEL 6.x have a xen-aware kernel built-in; therefore, there's no need to install xen-aware kernel package on theses versions.
- **install agent software**. This task is required for all supported versions.
- **enforce grub boot menu**. This task is required for all supported versions in order to boot the RN successfully, though the specific steps might vary for each version.

(RHEL 7.0) To enroll an agent-based Linux PN:

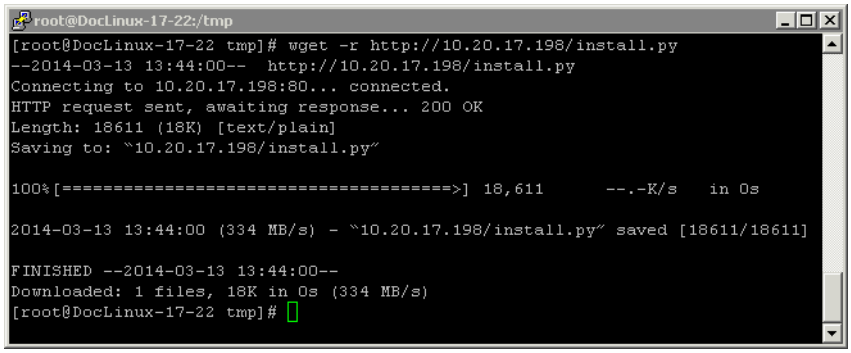
Use this procedure to enroll an agent-based Linux PN running RHEL 7.0.

1. Log on to the HA Appliance's onQ Portal as varadmin.
2. SSH to the to the server that you want to enroll as a PN, then log on as `root`.
3. Install the agent software:
   a. Launch the installer:

Note: By default, wget preserves the original files. Newly retrieved install.py is saved with extensions `.1`, `.2`, etc. Use the `-r` option (see wget man page) to overwrite the existing file.

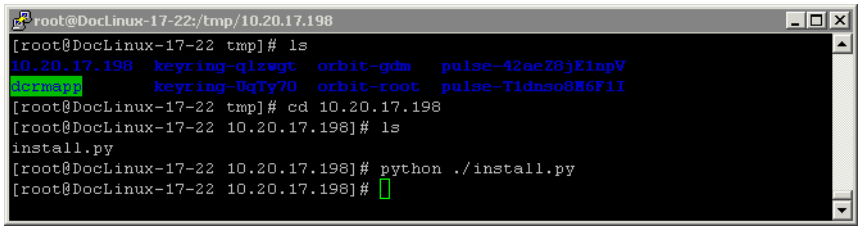
   b. From within a folder (`/tmp`) where you want to save the install script, run the following command:

```
# wget -r http://<onQ-IP-address>/install.py
```

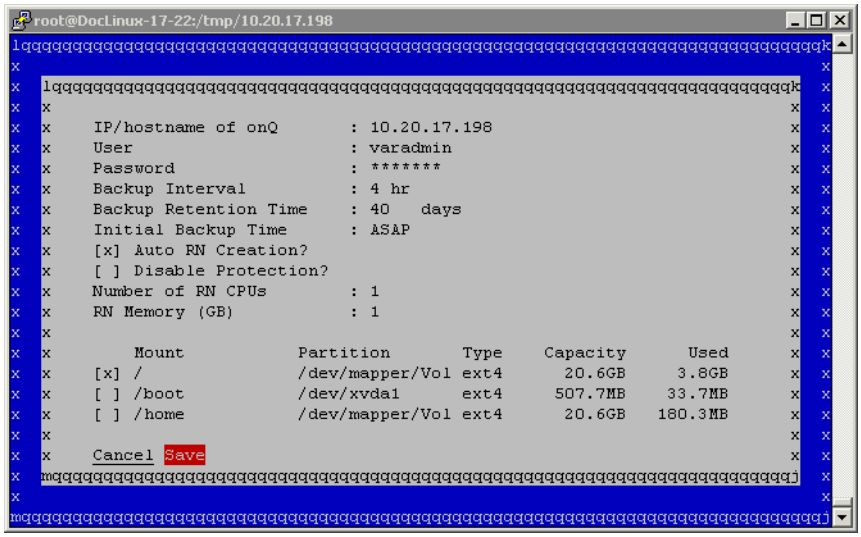


c. Start the installer:

```
# cd <onQ-IP-address>
# python ./install.py
```



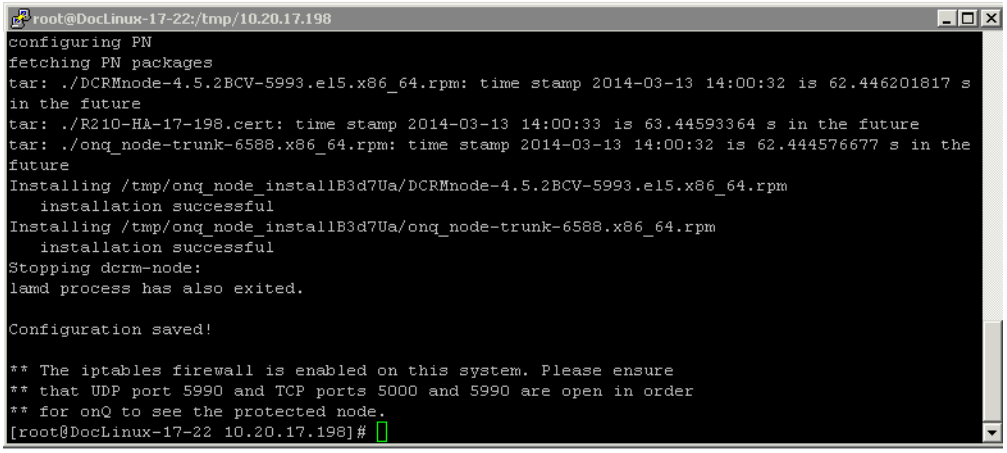
d. Type the credentials for either onQ varadmin or admin user.



e. (Optional) Specify all valid values, excluding file system type and capacity because these values aren't editable, in the node parameter fields provided and modify defaults as needed. (The

install utility lists all available volumes/partitions/mount points and the onQ portal enforces any file system requirements as outlined in Linux RN Limitations.) You can make these node changes at any point after enrollment via the **Protection Config** tab in the onQ Portal.

f. Select **Save** to update/install the client node package on the Linux node.



g. Wait! Do not press Enter. Was the installation successful? Use the error messages below to evaluate the screen output.

- If yes, exit the script.
- If no, press `Enter` to launch the installer UI again. Correct the problem, then repeat Step e through Step g.

`Completed Successfully`
 	
`The iptables firewall is enabled on this system....`
	Message appears in the shell, after the GUI curser exits. In addition, you'll be instructed to open up ports.

`Incorrect/invalid values entered`	Install utility stops if you type incorrect/invalid values. Correct the problem and **Save** again or **Cancel**.
`not authorized`	You either typed the credentials incorrectly or the user account does not have root privileges.

4. **Install the Netcat (`nc`) utility**, if not already installed. On your yum-enabled PN, run the following command, then verify that the package was installed. This utility reads and writes data across network connections using TCP or UDP. onQ depends on this utility for FLR restores.

```
# yum install nc
# which nc
nc is /usr/bin/nc
# rpm -qf /usr/bin/nc
nmap-ncat-6.40-4.el7.x86_64
```

5. **Create the grub boot menu**. In order to boot the RN successfully, the PN needs to prepare the init ram disk image with the required drivers and legacy grub boot menu.

   a. Verify OS version:

```
# cat /etc/redhat-release
Red Hat Enterprise Linux Server release 7.0
(Maipo)
```

   b. Make sure `ext2`/`ext3`/`ext4` file systems utilities are installed.

```
# rpm -qa | grep e2fsprogs
e2fsprogs-libs-1.42.9-4.el7.x86_64
e2fsprogs-1.42.9-4.el7.x86_64
```

   If not installed, do so now:

```
# yum install e2fsprogs
```

   c. Generate the init ram disk with xen drivers and `ext4` file system modules.

   Print the kernel release:

```
# uname -r
3.10.0-123.13.2.el7.x86_64
```

Where the `3.10.0-123.13.2.el7.x86_64` is the kernel release by default, change it to match PN's kernel release:

```
# cd /boot
# mkdir -p /boot/grub
# dracut --force --filesystems "ext4 ext3"  \
--add-drivers  "xen:vbd xen:vif" \
initramfs-3.10.0-123.13.2.el7xen.x86_64.img
```

d. Verify the legacy grub boot loader:

```
# vi /boot/grub/grub.conf.xvf5
```

Where the `3.10.0-123.13.2.el7.x86_64` is the kernel release by default, change `vmlinuz` and `initramfs` to match PN's kernel release. The kernel parameters are on a single line. Simply copy and paste from following screen.

Where the `root=UUID=855cd484-3852-4984-b568-ee0408c6b590`, the `855cd`... (UUID) is a temporary placeholder and will be replaced by read "`/`"'s UUID during the RN build. Do not make any changes to this parameter.

For example: The contents of `/boot/grub/grub.conf.xvf5`:

```
default=0
timeout=5
title onQ Red Hat Enterprise Linux (3.10.0-123.13.2.el7.x86_64)
root (hd0,0)
kernel /vmlinuz-3.10.0-123.13.2.el7.x86_64 ro root=UUID=855cd484-3852-
4984-b568-ee0408c6b590 plymouth.enable=0 console=hvc0 loglvl=all
cgroup_disable=memory sync_console console_to_ring earlyprintk=xen
nomodeset net.ifnames=0 biosdevname=0 LANG=en_US.UTF-8
initrd /initramfs-3.10.0-123.13.2el7xen.x86_64.img
```

From `/boot`, validate that `vmlinuz-3.10.0-123.13.2.el7.x86_64` and `initramfs-3.10.0-`

123.13.2.el7xen.x86_64.img exist in /boot folder as shown in the example below:

```
# ls /boot/vmlinuz-3.10.0-123.13.2.el7.x86_64
/boot/vmlinuz-3.10.0-123.13.2.el7.x86_64
# ls /boot/initramfs-3.10.0-
123.13.2.el7xen.x86_64.img
/boot/initramfs-3.10.0-123.13.2.el7xen.x86_64.img
```

6. **(Recommended) Disable the Network Manager service**, if installed, for self-tests to work correctly. This service is not useful in a server environment. Due to an RHEL 7.0 bug, turning off Network Manager can prevent NICs from showing up; if this occurs, re-enable the Network Manager service.

```
# systemctl disable NetworkManager.service
# systemctl stop NetworkManager.service
# systemctl status NetworkManager.service
```

7. **Wait for the RN to build, then perform a self-test**.

   Troubleshooting RN Build or Self-test Problems

   Mistakes with the grub boot menu enforcement can prevent the RN from booting. The following list represents the most common errors.

   - Kernel parameters are not on one single line. Some file editors wrap long parameters.
   - You have a typo in grub.con or grub.conf.xvf5 file name.
   - You have a typo in the kernel file name or the initramfs file name, or these files don’t exist.
   - There is a mismatch, on the boot menu, between the kernel versions and the initramfs version. If the kernel's version does not match the contents of initramfs, the RN won't boot.The system could have more than one kernel installed:

     7.0:

     vmlinuz-3.10.0-123.13.2.el7.x86_64

     should match

     initramfs-3.10.0-123.13.2.el7xen.x86_64.img

6.x:

`vmlinuz-2.6.32-279.el6.x86_64`

should match

`initramfs-2.6.32-279.el6.x86_64.img`

5.x:

`vmlinuz-2.6.18-371.el5xen`

should match

`initrd-2.6.18-371.el5xen.img.5`

To find the driver versions packed inside the init ram file system (`initramfs`) of the boot menu: Locate the initramfs and kernel name from the boot menu prepared for the RN (you'll find it under `/boot)`, then use the following command to peek the contents of initramfs. For example:

RHEL 6.x or 7.0:

```
# grep kernel /boot/grub/grub.conf.xvf5
kernel /vmlinuz-3.10.0-123.el7.x86_64  ro
root=UUID=9002ec24-fb30-4d16-8a78-b352a807e82b
plymouth.enable=0 console=hvc0 loglvl=all
cgroup_disable=memory sync_console
console_to_ring earlyprintk=xen nomodeset net.if-
names=0 biosdevname=0 LANG=en_US.UTF-8
# grep initrd /boot/grub/grub.conf.xvf5
initrd /initramfs-3.10.0-123.el7xen.x86_64.img
# lsinitrd /boot/initramsfs-3.10.0-
123.el7xen.x86_64.img|grep modules
rw-r--r--   1 root     root         1446 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.dep
-rw-r--r--   1 root     root         2450 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.dep.bin
-rw-r--r--   1 root     root           52 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.devname
-rw-r--r--   1 root     root        82512 Jun 30
2014 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.order
-rw-r--r--   1 root     root          165 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.softdep
-rw-r--r--   1 root     root        28132 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.symbols
-rw-r--r--   1 root     root        34833 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.symbols.bin
```

RHEL 5.x:

```
# grep kernel /boot/grub/grub.conf.xvf5
kernel /vmlinuz-2.6.18-371.el5xen ro
root=/dev/xvda1 rd_NO_LUKS rd_NO_MD rhgb crashker-
nel=auto rd_NO_LVM
# grep initrd /boot/grub/grub.conf.xvf5
initrd /initrd-2.6.18-371.el5xen.img.5
# zcat /tmp/initrd-2.6.18-371.el5xen.img.5|cpio
-t|grep -E "xen|ext"
16524 blocks
lib/ext3.ko
lib/xennet.ko
lib/xenblk.ko
```

8. **Log on to the PN and open the** following ports on the firewall in order for onQ to communicate with the PN.
   - UDP port `5990`
   - TCP ports `5000` and `5990`

Firewalld:

By default, RHEL 7.0 has introduced a new firewall service, a dynamic firewall daemon known as `firewalld`, instead of iptables service by default; however, the traditional iptables service is supported if installed. For details, see the Red Hat Linux 7 Security Guide. If you choose to disable firewalld, there is no need to configure firewalld firewall rules: simply skip this procedure.

`firewalld` daemon and service and iptables service are using iptables commands to configure the netfilter in the kernel to separate and filter the network traffic. `firewalld` stores it in various XML files in `/usr/lib/firewalld/` and `/etc/firewalld/`.

The firewalld firewall defines how *networks zones* can be used to separate networks into different zones. Based on the level of trust, you can decide to place devices and traffic within a particular network zone. Each mutable network zone can have a different combination of firewall rules.

a. Verify that firewalld is in a running state.

b. Check the service status:

```
[root@RHEL70x64-17-167 services]# systemctl status firewalld.service
firewalld.service - firewalld - dynamic firewall daemon
Loaded: loaded (/usr/lib/systemd/system/fire-walld.service; disabled)
Active: inactive (dead)
```

c. Enable the service, if not already enabled:

```
[root@RHEL70x64-17-167 services]# systemctl enable firewalld
ln -s '/usr/lib/systemd/system/firewalld.service' '/etc/systemd/system/dbusorg.edorapro-ject.FirewallD1.service'
ln -s '/usr/lib/systemd/system/firewalld.service' '/etc/systemd/system
/basic.target.wants/firewalld.service'
```

d. Start the service, if not already running:

```
[root@RHEL70x64-17-167 services]# systemctl start firewalld
```

e. On the PN, find the network interface that is used to communicate with onQ. In this example, that NIC is `ens32`.

```
[root@RHEL70x64-17-167 services]# ifconfig
ens32: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>
mtu 1500
inet 10.20.17.167 netmask 255.255.248.0 broadcast
10.20.23.255
inet6 fe80::250:56ff:fe9d:2121 prefixlen 64 sco-
peid 0x20<link>
ether 00:50:56:9d:21:21 txqueuelen 1000 (Ethernet)
RX packets 7115713 bytes 476287831 (454.2 MiB)
RX errors 0 dropped 149791 overruns 0 frame 0
TX packets 924966 bytes 1305413839 (1.2 GiB)
TX errors 0 dropped 0 overruns 0 carrier 0 colli-
sions 0

lo: flags=73<UP,LOOPBACK,RUNNING> mtu 65536
inet 127.0.0.1 netmask 255.0.0.0
inet6 ::1 prefixlen 128 scopeid 0x10<host>
loop txqueuelen 0 (Local Loopback)
RX packets 10 bytes 980 (980.0 B)
RX errors 0 dropped 0 overruns 0 frame 0
TX packets 10 bytes 980 (980.0 B)
TX errors 0 dropped 0 overruns 0 carrier 0 colli-
sions 0
```

f. For the network interface that you identified above, find the network interface's network zone. In this example, the network zone is `work`.

```
[root@RHEL70x64-17-167 services]# firewall-cmd --get-
zone-of-interface=ens32
work
```

Determine your default zone. In the following example, the default zone is `Public`.

```
[root@RHEL70x64-17-167 services]# firewall-cmd --get-default-zone
public
```

g. Associate the zone(s) with the following firewall rules. The same rules can be applied to many zones as needed. In the following example, dcrm-node service is associated with `work` zone for `ens32`. The `dcrm-node.xml` is located at `/usr/lib/firewalld/services`.

```
[root@RHEL70x64-17-167 services]# firewall-cmd --add-service=dcrm-node --permanent --zone=work
success
```

h. Activate the latest firewall rules:

```
[root@RHEL70x64-17-167 services]# firewall-cmd --reload
success
```

Now the PN can communicate with onQ.

i. Set up the rule for RN on the PN site.

The RN will be equipped with `eth0` interface, so apply the rules to `eth0` interface's zone if different from PN's zone. The PN might not have `eth0` interface; in such a case, the RN's `eth0` will be in the default zone.

Find `eth0` network interface's network zone. In this example, it is *not* set:

```
[root@RHEL70x64-17-167 services]# firewall-cmd --get-zone-of-interface=eth0
no zone
```

Determine your default zone. In this example default zone is `Public`. Since `eth0` has no zone dcm-node is associated with `Public` zone:

```
[root@RHEL70x64-17-167 services]# firewall-cmd --get-default-zone
public
```

j. Associate the zone(s) with the following firewall rules. The same rules can be applied to many zones as needed:

```
[root@RHEL70x64-17-167 services]# firewall-cmd --add-service=dcrm-node --permanent --zone=public
success
```

k. Active the latest firewall rules:

```
[root@RHEL70x64-17-167 services]# firewall-cmd --reload
success
```

Now the RN can communicate with onQ, mainly for self-tests.

I. Confirm the firewall rules. The public zone and work zone have TCP ports (5000/5990) and UDP port 5990 opened in this case.

```
[root@RHEL70x64-17-167 services]# iptables -L -n
Chain IN_public (2 references)
target prot opt source destination
IN_public_log all -- 0.0.0.0/0 0.0.0.0/0
IN_public_deny all -- 0.0.0.0/0 0.0.0.0/0
IN_public_allow all -- 0.0.0.0/0 0.0.0.0/0
Chain IN_public_allow (1 references)
target prot opt source destination
ACCEPT tcp -- 0.0.0.0/0 0.0.0.0/0 tcp dpt:5990 ct-
state NEW
ACCEPT tcp -- 0.0.0.0/0 0.0.0.0/0 tcp dpt:5000 ct-
state NEW
ACCEPT udp -- 0.0.0.0/0 0.0.0.0/0 udp dpt:5990 ct-
state NEW
ACCEPT tcp -- 0.0.0.0/0 0.0.0.0/0 tcp dpt:22 ct-
state NEW
Chain IN_work (0 references)
target prot opt source destination
IN_work_log all -- 0.0.0.0/0 0.0.0.0/0
IN_work_deny all -- 0.0.0.0/0 0.0.0.0/0
IN_work_allow all -- 0.0.0.0/0 0.0.0.0/0
Chain IN_work_allow (1 references)
target prot opt source destination
ACCEPT udp -- 0.0.0.0/0 0.0.0.0/0 udp dpt:631 ct-
state NEW
ACCEPT tcp -- 0.0.0.0/0 0.0.0.0/0 tcp dpt:5990 ct-
state NEW
ACCEPT tcp -- 0.0.0.0/0 0.0.0.0/0 tcp dpt:5000 ct-
state NEW
ACCEPT udp -- 0.0.0.0/0 0.0.0.0/0 udp dpt:5990 ct-
state NEW
ACCEPT tcp -- 0.0.0.0/0 0.0.0.0/0 tcp dpt:22 ct-
state NEW
```

iptables:

a. Verify that udp 5990 and tcp 5000 and 5990 ports are open and above the `REJECT` line:

```
# iptables -L -n | grep -E "5900|5000"
ACCEPT     udp  --  0.0.0.0/0           0.0.0.0/0
udp dpt:5990
ACCEPT     tcp  --  0.0.0.0/0           0.0.0.0/0
state NEW tcp dpt:5000
ACCEPT     tcp  --  0.0.0.0/0           0.0.0.0/0
state NEW tcp dpt:5990
REJECT all -- 0.0.0.0/0 0.0.0.0/0 reject-with
icmp-host-prohibited
```

b. If these ports are not open, open them.

Find the input chain name, which is reported in the line that appears after `INPUT`:

```
# iptables -L -line numbers
Chain INPUT (policy ACCEPT)
num target prot opt source destination
1 RH-Firewall-1-INPUT all -- anywhere anywhere
...
10 REJECT all -- anywhere anywhere reject-with
icmp-host-prohibited
```

Using the line number of the `REJECT` line and the chain name (line `10` and line `1`, respectively, in the step above), insert the onQ ports:

```
# iptables -I RH-Firewall-1-INPUT 10 -m state --
state NEW -p tcp --dport 5990 -j ACCEPT
# iptables -I RH-Firewall-1-INPUT 10 -m state --
state NEW -p tcp --dport 5000 -j ACCEPT
# iptables -I RH-Firewall-1-INPUT 10 -p udp --
dport 5990 -j ACCEPT
```

c. Save and restart iptables.

Afterward, verify that the ports are open and above the `REJECT` line as outlined earlier.

```
# service iptables save
Saving firewall rules to /etc/sysconfig/iptables:
[ OK ]
# service iptables restart
Flushing firewall rules: [ OK ]
Setting chains to policy ACCEPT: filter [ OK ]
Unloading iptables modules: [ OK ]
Applying iptables firewall rules: [ OK ]
Loading additional iptables modules:
ip_conntrack_netbios_n [ OK ]
```

9. **(Oracle database) Install `RMAN` scripts**.

   If the PN has an Oracle database, install the `RMAN` scripts so that onQ can execute a hot backup of your database as outlined in Back up and restore Oracle 11g database on Linux.

(RHEL 6.x) To enroll an agent-based Linux PN:

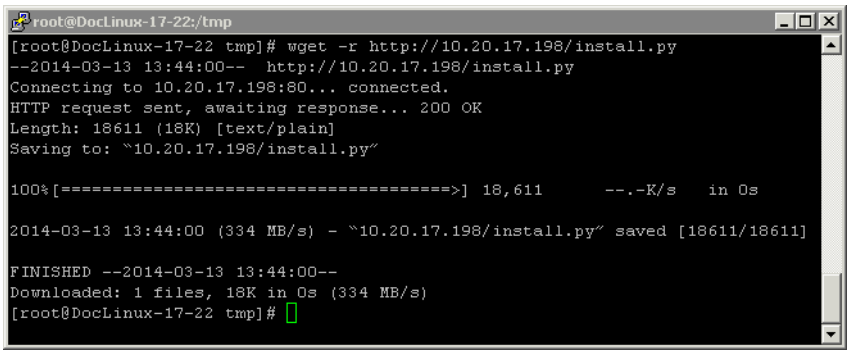
Use this procedure to enroll an agent-based Linux PN running RHEL 6.x.

1. Log on to the HA Appliance's onQ Portal as varadmin.
2. SSH to the to the server that you want to enroll as a PN, then log on as `root`.
3. Install the agent software:
   a. Launch the installer:

Note: By default, wget preserves the original files. Newly retrieved install.py is saved with extensions `.1`, `.2`, etc. Use the `-r` option (see wget man page) to overwrite the existing file.

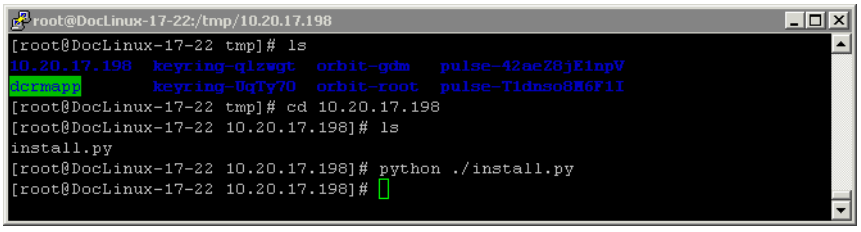
b. From within a folder (`/tmp`) where you want to save the install script, run the following command:

```
# wget -r http://<onQ-IP-address>/install.py
```

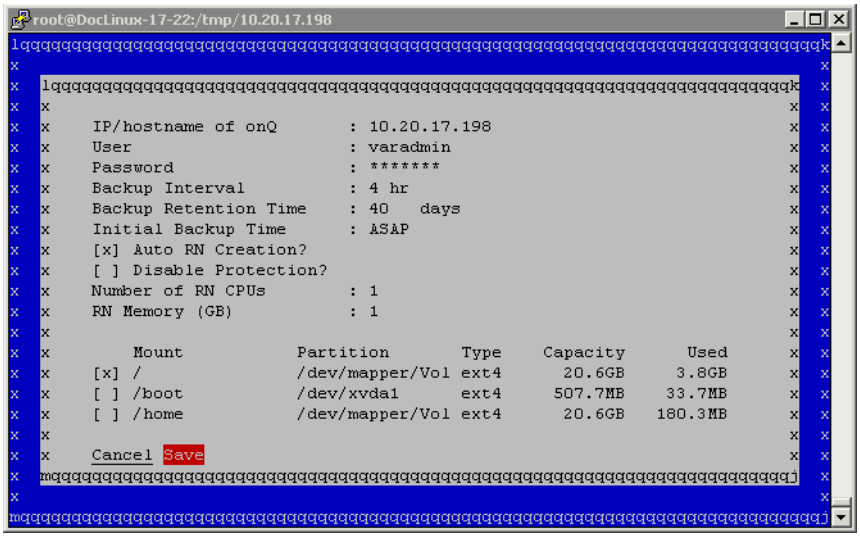


c. Start the installer:

```
# cd <onQ-IP-address>
# python ./install.py
```

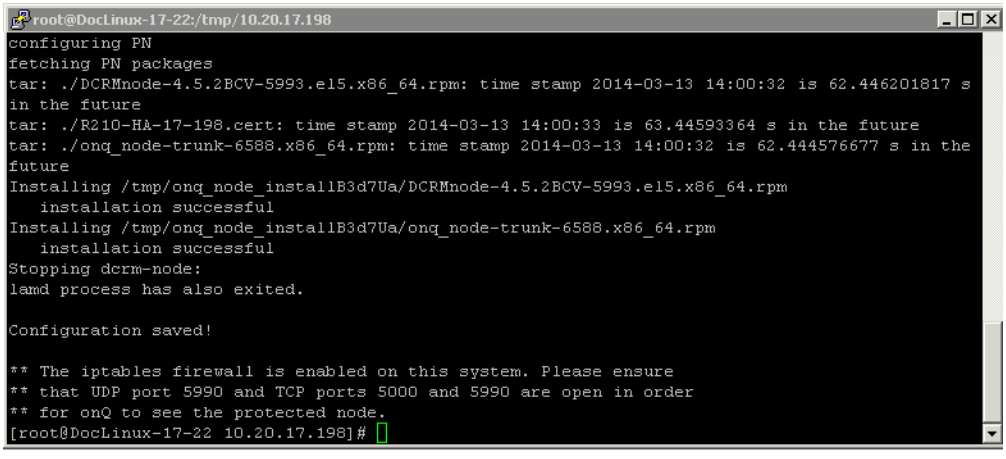


d. Type the credentials for either onQ varadmin or admin user.



e. (Optional) Specify all valid values, excluding file system type and capacity because these values aren't editable, in the node parameter fields provided and modify defaults as needed. (The install utility lists all available volumes/partitions/mount points and the onQ portal enforces any file system requirements as outlined in Linux RN Limitations.) You can make these node changes at any point after enrollment via the **Protection Config** tab in the onQ Portal.

f. Select **Save** to update/install the client node package on the Linux node.



g. Wait! Do not press Enter. Was the installation successful? Use the error messages below to evaluate the screen output.

- If yes, exit the script.
- If no, press `Enter` to launch the installer UI again. Correct the problem, then repeat Step e through Step g.

`Completed Successfully`
 	
`The iptables firewall is enabled on this system....`
	Message appears in the shell, after the GUI curser exits. In addition, you'll be instructed to open up ports.

`Incorrect/invalid values entered`	Install utility stops if you type incorrect/invalid values. Correct the problem and **Save** again or **Cancel**.
`not authorized`	You either typed the credentials incorrectly or the user account does not have root privileges.

4. **Install the Netcat (`nc`) utility**, if not already installed. On your yum-enabled PN, run the following command, then verify that the package was installed. This utility reads and writes data across network connections using TCP or UDP. onQ depends on this utility for FLR restores.

```
# yum install nc
# which nc
nc is /usr/bin/nc
# rpm -qf /usr/bin/nc
nc-1.84-22.el6.x86_64
```

5. **Wait for the RN to build, then perform a self-test**.

   Troubleshooting RN Build or Self-test Problems

   Mistakes with the grub boot menu enforcement can prevent the RN from booting. The following list represents the most common errors.
   - Kernel parameters are not on one single line. Some file editors wrap long parameters.
   - You have a typo in `grub.con` or `grub.conf.xvf5` file name.
   - You have a typo in the kernel file name or the initramfs file name, or these files don't exist.

- There is a mismatch, on the boot menu, between the kernel versions and the initramfs version. If the kernel's version does not match the contents of initramfs, the RN won't boot.The system could have more than one kernel installed:

  7.0:

  `vmlinuz-3.10.0-123.13.2.el7.x86_64`

  should match

  `initramfs-3.10.0-123.13.2.el7xen.x86_64.img`

  6.x:

  `vmlinuz-2.6.32-279.el6.x86_64`

  should match

  `initramfs-2.6.32-279.el6.x86_64.img`

  5.x:

  `vmlinuz-2.6.18-371.el5xen`

  should match

  `initrd-2.6.18-371.el5xen.img.5`

  To find the driver versions packed inside the init ram file system (`initramfs`) of the boot menu: Locate the initramfs and kernel name from the boot menu prepared for the RN (you'll find it under `/boot)`, then use the following command to peek the contents of initramfs. For example:

RHEL 6.x or 7.0:

```
# grep kernel /boot/grub/grub.conf.xvf5
kernel /vmlinuz-3.10.0-123.el7.x86_64  ro
root=UUID=9002ec24-fb30-4d16-8a78-b352a807e82b
plymouth.enable=0 console=hvc0 loglvl=all
cgroup_disable=memory sync_console
console_to_ring earlyprintk=xen nomodeset net.if-
names=0 biosdevname=0 LANG=en_US.UTF-8
# grep initrd /boot/grub/grub.conf.xvf5
initrd /initramfs-3.10.0-123.el7xen.x86_64.img
# lsinitrd /boot/initramsfs-3.10.0-
123.el7xen.x86_64.img|grep modules
rw-r--r--   1 root     root         1446 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.dep
-rw-r--r--   1 root     root         2450 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.dep.bin
-rw-r--r--   1 root     root           52 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.devname
-rw-r--r--   1 root     root        82512 Jun 30
2014 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.order
-rw-r--r--   1 root     root          165 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.softdep
-rw-r--r--   1 root     root        28132 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.symbols
-rw-r--r--   1 root     root        34833 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.symbols.bin
```

RHEL 5.x:

```
# grep kernel /boot/grub/grub.conf.xvf5
kernel /vmlinuz-2.6.18-371.el5xen ro
root=/dev/xvda1 rd_NO_LUKS rd_NO_MD rhgb crashker-
nel=auto rd_NO_LVM
# grep initrd /boot/grub/grub.conf.xvf5
initrd /initrd-2.6.18-371.el5xen.img.5
# zcat /tmp/initrd-2.6.18-371.el5xen.img.5|cpio
-t|grep -E “xen|ext”
16524 blocks
lib/ext3.ko
lib/xennet.ko
lib/xenblk.ko
```

6. **(Recommended) Disable the Network Manager service**, if installed, for self-tests to work correctly. This service is not useful in a server environment.

```
# service NetworkManager stop
# service NetworkManager status
# chkconfig NetworkManager off
# chkconfig --list NetworkManager
NetworkManager  0:off   1:off   2:off   3:off
4:off   5:off   6:off
```

7. **Log on to the PN and open the** following ports on the firewall in order for onQ to communicate with the PN.
   - UDP port `5990`
   - TCP ports `5000` and `5990`

iptables:

a. Verify that udp 5990 and tcp 5000 and 5990 ports are open and above the `REJECT` line:

```
# iptables -L -n | grep -E "5900|5000"
ACCEPT     udp  --  0.0.0.0/0           0.0.0.0/0
udp dpt:5990
ACCEPT     tcp  --  0.0.0.0/0           0.0.0.0/0
state NEW tcp dpt:5000
ACCEPT     tcp  --  0.0.0.0/0           0.0.0.0/0
state NEW tcp dpt:5990
REJECT all -- 0.0.0.0/0 0.0.0.0/0 reject-with
icmp-host-prohibited
```

b. If these ports are not open, open them.

Find the input chain name, which is reported in the line that appears after `INPUT`:

```
# iptables -L -line numbers
Chain INPUT (policy ACCEPT)
num target prot opt source destination
1 RH-Firewall-1-INPUT all -- anywhere anywhere
...
10 REJECT all -- anywhere anywhere reject-with
icmp-host-prohibited
```

Using the line number of the `REJECT` line and the chain name (line `10` and line `1`, respectively, in the step above), insert the onQ ports:

```
# iptables -I RH-Firewall-1-INPUT 10 -m state --
state NEW -p tcp --dport 5990 -j ACCEPT
# iptables -I RH-Firewall-1-INPUT 10 -m state --
state NEW -p tcp --dport 5000 -j ACCEPT
# iptables -I RH-Firewall-1-INPUT 10 -p udp --
dport 5990 -j ACCEPT
```

c. Save and restart iptables.

Afterward, verify that the ports are open and above the `REJECT` line as outlined earlier.

```
# service iptables save
Saving firewall rules to /etc/sysconfig/iptables:
[ OK ]
# service iptables restart
Flushing firewall rules: [ OK ]
Setting chains to policy ACCEPT: filter [ OK ]
Unloading iptables modules: [ OK ]
Applying iptables firewall rules: [ OK ]
Loading additional iptables modules:
ip_conntrack_netbios_n [ OK ]
```

8. **(Oracle database) Install `RMAN` scripts**.

   If the PN has an Oracle database, install the `RMAN` scripts so that onQ can execute a hot backup of your database as outlined in Back up and restore Oracle 11g database on Linux.

(RHEL 5.x) To enroll an agent-based Linux PN:

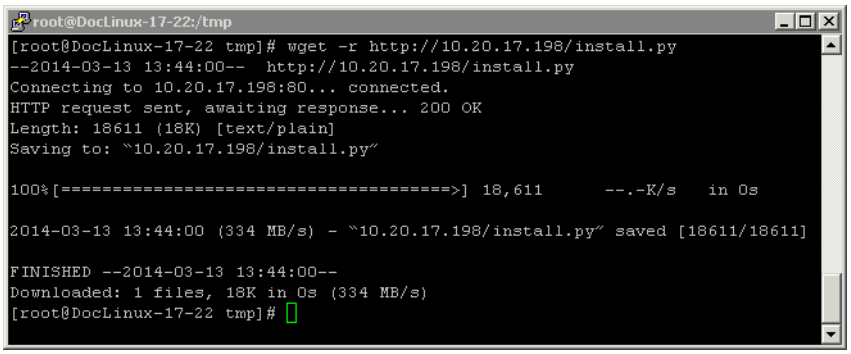
Use this procedure to enroll an agent-based Linux PN running RHEL 5.x.

1. Log on to the HA Appliance's onQ Portal as varadmin.
2. SSH to the to the server that you want to enroll as a PN, then log on as `root`.
3. Install the agent software:
   a. Launch the installer:

Note: By default, wget preserves the original files. Newly retrieved install.py is saved with extensions `.1`, `.2`, etc. Use the `-r` option (see wget man page) to overwrite the existing file.

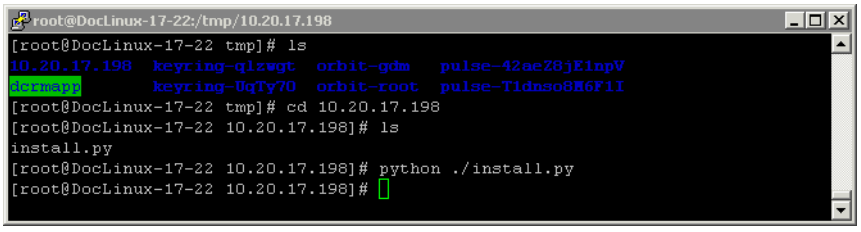
b. From within a folder (`/tmp`) where you want to save the install script, run the following command:

```
# wget -r http://<onQ-IP-address>/install.py
```

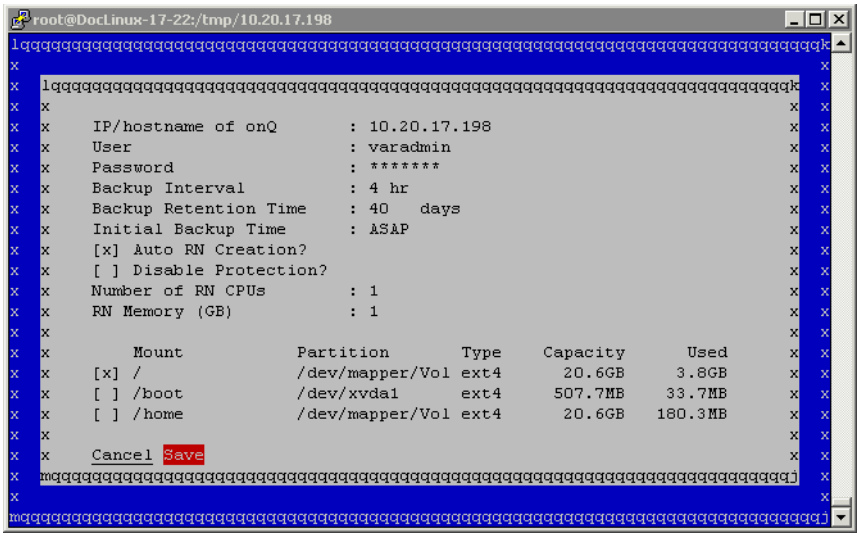


c. Start the installer:

```
# cd <onQ-IP-address>
# python ./install.py
```

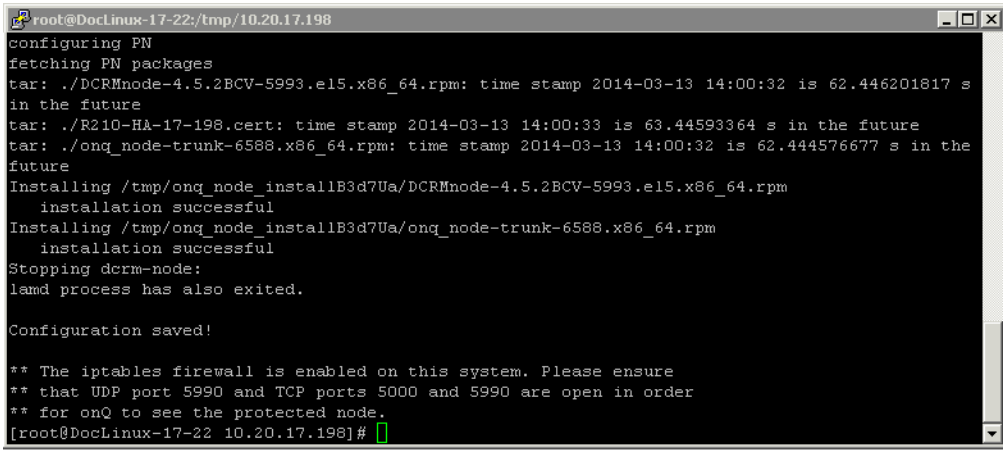


d. Type the credentials for either onQ varadmin or admin user.



e. (Optional) Specify all valid values, excluding file system type and capacity because these values aren't editable, in the node parameter fields provided and modify defaults as needed. (The install utility lists all available volumes/partitions/mount points and the onQ portal enforces any file system requirements as outlined in Linux RN Limitations.) You can make these node changes at any point after enrollment via the **Protection Config** tab in the onQ Portal.

f. Select **Save** to update/install the client node package on the Linux node.



g. Wait! Do not press Enter. Was the installation successful? Use the error messages below to evaluate the screen output.

- If yes, exit the script.
- If no, press `Enter` to launch the installer UI again. Correct the problem, then repeat Step e through Step g.

`Completed Successfully`
 	
`The iptables firewall is enabled on this system....`
	Message appears in the shell, after the GUI curser exits. In addition, you'll be instructed to open up ports.

`Incorrect/invalid values entered`	Install utility stops if you type incorrect/invalid values. Correct the problem and **Save** again or **Cancel**.
`not authorized`	You either typed the credentials incorrectly or the user account does not have root privileges.

4. Install kernel-xen RPM package.

5. Copy and modify `/boot/grub/menu.lst`:

```
default=0
timeout=5
hiddenmenu
title Red Hat Enterprise Linux Server by Quorum onQ (2.6.18-371.el5xen)
root (hd0,0)
kernel /vmlinuz-2.6.18-371.el5xen ro root=/dev/xvda1 rd_NO_LUKS rd_NO_MD rhgb crashkernel=auto rd_NO_LVM
initrd /initrd-2.6.18-371.el5xen.img.5
```

6. **Wait for the RN to build, then perform a self-test**.

   Troubleshooting RN Build or Self-test Problems

   Mistakes with the grub boot menu enforcement can prevent the RN from booting. The following list represents the most common errors.

   - Kernel parameters are not on one single line. Some file editors wrap long parameters.
   - You have a typo in `grub.con` or `grub.conf.xvf5` file name.

- You have a typo in the kernel file name or the initramfs file name, or these files don’t exist.
- There is a mismatch, on the boot menu, between the kernel versions and the initramfs version. If the kernel's version does not match the contents of initramfs, the RN won't boot.The system could have more than one kernel installed:

  7.0:

  `vmlinuz-3.10.0-123.13.2.el7.x86_64`

  should match

  `initramfs-3.10.0-123.13.2.el7xen.x86_64.img`

  6.x:

  `vmlinuz-2.6.32-279.el6.x86_64`

  should match

  `initramfs-2.6.32-279.el6.x86_64.img`

  5.x:

  `vmlinuz-2.6.18-371.el5xen`

  should match

  `initrd-2.6.18-371.el5xen.img.5`

  To find the driver versions packed inside the init ram file system (`initramfs`) of the boot menu: Locate the initramfs and kernel name from the boot menu prepared for the RN (you’ll find it under `/boot)`, then use the following command to peek the contents of initramfs. For example:

RHEL 6.x or 7.0:

```
# grep kernel /boot/grub/grub.conf.xvf5
kernel /vmlinuz-3.10.0-123.el7.x86_64  ro
root=UUID=9002ec24-fb30-4d16-8a78-b352a807e82b
plymouth.enable=0 console=hvc0 loglvl=all
cgroup_disable=memory sync_console
console_to_ring earlyprintk=xen nomodeset net.if-
names=0 biosdevname=0 LANG=en_US.UTF-8
# grep initrd /boot/grub/grub.conf.xvf5
initrd /initramfs-3.10.0-123.el7xen.x86_64.img
# lsinitrd /boot/initramsfs-3.10.0-
123.el7xen.x86_64.img|grep modules
rw-r--r--   1 root     root         1446 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.dep
-rw-r--r--   1 root     root         2450 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.dep.bin
-rw-r--r--   1 root     root           52 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.devname
-rw-r--r--   1 root     root        82512 Jun 30
2014 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.order
-rw-r--r--   1 root     root          165 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.softdep
-rw-r--r--   1 root     root        28132 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.symbols
-rw-r--r--   1 root     root        34833 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.symbols.bin
```

RHEL 5.x:

```
# grep kernel /boot/grub/grub.conf.xvf5
kernel /vmlinuz-2.6.18-371.el5xen ro
root=/dev/xvda1 rd_NO_LUKS rd_NO_MD rhgb crashker-
nel=auto rd_NO_LVM
# grep initrd /boot/grub/grub.conf.xvf5
initrd /initrd-2.6.18-371.el5xen.img.5
# zcat /tmp/initrd-2.6.18-371.el5xen.img.5|cpio
-t|grep -E “xen|ext”
16524 blocks
lib/ext3.ko
lib/xennet.ko
lib/xenblk.ko
```

7. **(Recommended) Disable the Network Manager service and Kudzu services**, if installed, for self-tests to work correctly. This service is not useful in a server environment.

```
# service NetworkManager stop
# service NetworkManager status
# chkconfig NetworkManager off
# chkconfig --list NetworkManager
NetworkManager  0:off   1:off   2:off   3:off
4:off   5:off   6:off
# service kudzu stop
# service kudzu status
# chkconfig kudzu off
# chkconfig --list kudzu
kudzu         0:off  1:off  2:off  3:off  4:off
5:off   6:off
```

8. **Log on to the PN and open the** following ports on the firewall in order for onQ to communicate with the PN.
   - UDP port `5990`
   - TCP ports `5000` and `5990`

iptables:

a. Verify that udp 5990 and tcp 5000 and 5990 ports are open and above the `REJECT` line:

```
# iptables -L -n | grep -E "5900|5000"
ACCEPT     udp  --  0.0.0.0/0            0.0.0.0/0
udp dpt:5990
ACCEPT     tcp  --  0.0.0.0/0            0.0.0.0/0
state NEW tcp dpt:5000
ACCEPT     tcp  --  0.0.0.0/0            0.0.0.0/0
state NEW tcp dpt:5990
REJECT all -- 0.0.0.0/0 0.0.0.0/0 reject-with
icmp-host-prohibited
```

b. If these ports are not open, open them.

Find the input chain name, which is reported in the line that appears after `INPUT`:

```
# iptables -L -line numbers
Chain INPUT (policy ACCEPT)
num target prot opt source destination
1 RH-Firewall-1-INPUT all -- anywhere anywhere
...
10 REJECT all -- anywhere anywhere reject-with
icmp-host-prohibited
```

Using the line number of the `REJECT` line and the chain name (line `10` and line `1`, respectively, in the step above), insert the onQ ports:

```
# iptables -I RH-Firewall-1-INPUT 10 -m state --
state NEW -p tcp --dport 5990 -j ACCEPT
# iptables -I RH-Firewall-1-INPUT 10 -m state --
state NEW -p tcp --dport 5000 -j ACCEPT
# iptables -I RH-Firewall-1-INPUT 10 -p udp --
dport 5990 -j ACCEPT
```

c. Save and restart iptables.

Afterward, verify that the ports are open and above the `REJECT` line as outlined earlier.

```
# service iptables save
Saving firewall rules to /etc/sysconfig/iptables:
[ OK ]
# service iptables restart
Flushing firewall rules: [ OK ]
Setting chains to policy ACCEPT: filter [ OK ]
Unloading iptables modules: [ OK ]
Applying iptables firewall rules: [ OK ]
Loading additional iptables modules:
ip_conntrack_netbios_n [ OK ]
```

9. **(Oracle database) Install `RMAN` scripts**.

   If the PN has an Oracle database, install the `RMAN` scripts so that onQ can execute a hot backup of your database as outlined in Back up and restore Oracle 11g database on Linux.

(Optional) To enforce a network setting:

Unlike with Windows RNs, Linux RNs cannot be assigned to networks as outlined in Assign RNs to Networks. However, you can still enforce network settings for RN self-tests.

The Linux RN build should have a working NIC. If you have special need to set the RN's NIC to a different IP, you can place a file, `/opt/quorum/bin/xvf.dat`, to generate `/etc/sysconfig/network-scripts/ifcfg-eth0`.

For example:

```
# cat /opt/quorum/bin/xvf.dat
@@XV_PN_IP
10.20.16.74
@@XV_PN_MASK
255.255.248.0
@@XV_PN_GW
10.20.16.1
```

Related Topics

(Agent-based Linux PNs) Update node software
(Agent-based Windows PNs) Enroll protected nodes
(Agent-less Linux/Windows PNs) Enroll protected nodes
(Agent-based PNs) Connection Problems

5.3 (Agent-less Linux/Windows PNs) Enroll protected nodes

onQ supports both agent-based and agent-less PNs. To protect agent-less PNs, you indirectly enroll PNs by enrolling the ESX/ESXi server itself. onQ's VMware Real Agentless protection provides the following benefits:

- No PN CPU utilization.
- Easier deployment in a VMware environment.
- No need to test or revalidate servers due to the introduction of new software.
- No need to worry about voiding vendor contracts or support due to installing agent-based software.

However, before you deploy agent-less PNs, consider the tradeoffs outlined in "Agent-less PN Enrollment Limitations" in onQ Release Notes.

VMware-hosted nodes, whether Windows or Linux, do not require that an agent (onQ Service) be installed on the PN. Instead, two *PN proxies* will be automatically installed on the ESX/ESXi server (aka *Proxy Host*) when you enroll that ESX/ESXi server: one Linux PN proxy (`LPY_`<onQhostname>`_`<ESXhostname>) and one Windows PN proxy (`WPY_`<onQhostname>`_`<ESXhostname>) for each onQ. All PNs on a given ESX/ESXi server and having the same operating system (Linux or Windows) will share a common PN proxy.

Using these proxies, the onQ's backup utility works with VMware Snapshot to back up the PNs to the HA Appliance.

Note: If you need to rename the onQ host name, you will need to perform the steps outlined in Modify Appliance's network settings.

(Windows) To enroll an agent-less Windows PN:

Use this procedure to enroll an agent-less Windows PN.

1. **Install the PN proxies**:
   a. Enable `ssh` on the ESXi/ESX server.

b. Reserve two static IP address, one for the Windows PN proxy and one for the Linux PN proxy.

c. Log on to the HA Appliance's onQ Portal.

d. Click the **PROTECTION CONFIG** tab > double-plus button (**++**).

The **Add Protected Nodes via Host** dialog appears.

e. Provide the credentials and either the hostname or the IP address for the ESXi/ESX host, then **GET LIST**.

The onQ Portal queries the virtual machines hosted by the ESXi/ESX server, then displays the inventory.

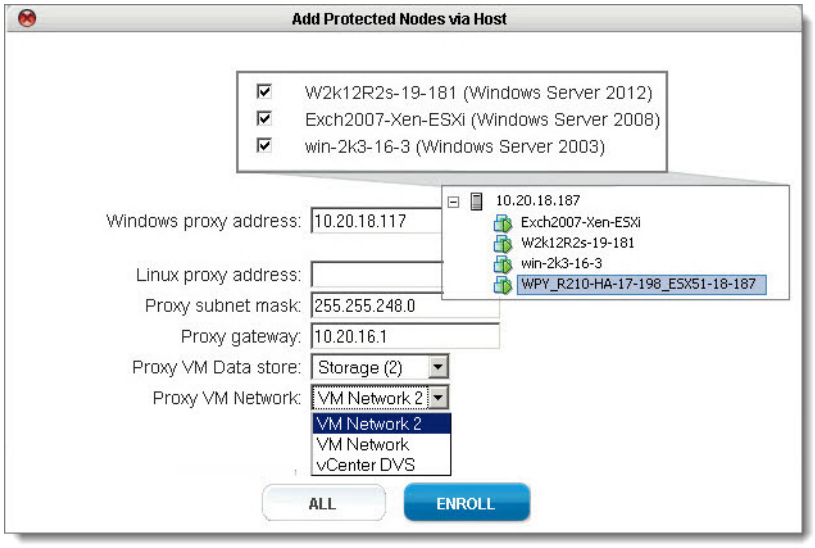
f. In the virtual machine list, select the check boxes (or **ALL** button) for the PNs that you want to enroll.

If you do not see a virtual machine does not appear in the list, it's likely that the virtual machine does not have VMware Tools installed.

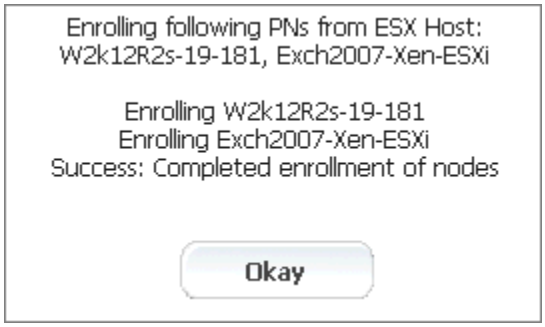
Provide the following information to install the proxies:

- **Windows proxy address**. The static IP address that you reserved for the Windows PN proxy.
- **Linux proxy address**. The static IP address that you reserved for the Linux PN proxy.
- **Proxy subnet mask**. The PN proxy's subnet mask.
- **Proxy gateway**. The PN proxy's gateway.
- **Proxy VM Data store**. The directory in which the vm's files are stored.
- **Proxy VM Network**. Your network configuration can include a vCenter-level Distributed Virtual Switches (DVS) or an ESXi host-level virtual switch. When you enroll an ESXi host, all available networks that are visible to that host are listed and available for selection.

In the following example, three Windows PN are being enrolled, so only one static IP address (for the Windows PN proxy `WPY_`<onQhostname>`_`<ESXhostname>) is needed:



Click the **ENROLL** button, then **OKAY**.



The PNs that you selected appear in the Protected Nodes list; however, they are `unverified` as evidenced by the status in the **PROTECTED NODES** page > **Protection Disabled** column.

g. From the ESXi/ESX host, power on the PN proxies.

h. Activate (aka *verify*) the PNs so that onQ can protect them.

Go to **PROTECTION CONFIG** tab, select the PN, then click the **MODIFY** button.

Specify values for the node parameters, then **SAVE**. The act of saving a PN's configuration instructs onQ to activate that PN's configuration.

The onQ Portal groups the PNs by ESXi server. If you do not want these PNs in such a group, clear the default **Group Name** field to remove the PNs from this group and place them in the shared pool.

2. **Log on to the PN and open the** following ports on the firewall in order for onQ to communicate with the PN.
   - UDP port `5990`
   - TCP ports `5000` and `5990`
3. (Oracle database) Install `RMAN` scripts.

   If the PN has an Oracle database, install the `RMAN` scripts so that onQ can execute a hot backup of your database as outlined in Back up and restore Oracle 10g+ database on Windows.

(RHEL 7.0) To enroll an agent-less Linux PN:

Use this procedure to enroll an agent-less Linux PN running RHEL 7.0.

1. **Install the PN proxies**:
   a. Enable `ssh` on the ESXi/ESX server.
   b. Reserve two static IP address, one for the Windows PN proxy and one for the Linux PN proxy.
   c. Log on to the HA Appliance's onQ Portal.
   d. Click the **PROTECTION CONFIG** tab > double-plus button (**++**).

      The **Add Protected Nodes via Host** dialog appears.
   e. Provide the credentials and either the hostname or the IP address for the ESXi/ESX host, then **GET LIST**.

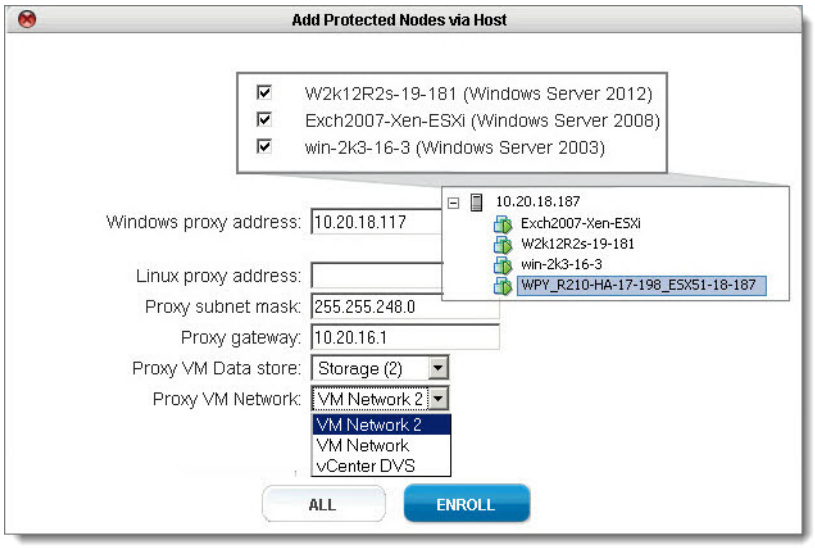
      The onQ Portal queries the virtual machines hosted by the ESXi/ESX server, then displays the inventory.
   f. In the virtual machine list, select the check boxes (or **ALL** button) for the PNs that you want to enroll.

      If you do not see a virtual machine does not appear in the list, it's likely that the virtual machine does not have VMware Tools installed.

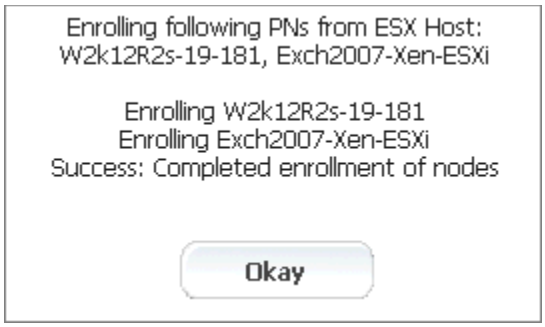
      Provide the following information to install the proxies:

- **Windows proxy address**. The static IP address that you reserved for the Windows PN proxy.
- **Linux proxy address**. The static IP address that you reserved for the Linux PN proxy.
- **Proxy subnet mask**. The PN proxy's subnet mask.
- **Proxy gateway**. The PN proxy's gateway.
- **Proxy VM Data store**. The directory in which the vm's files are stored.
- **Proxy VM Network**. Your network configuration can include a vCenter-level Distributed Virtual Switches (DVS) or an ESXi host-level virtual switch. When you enroll an ESXi host, all available networks that are visible to that host are listed and available for selection.

In the following example, three Windows PN are being enrolled, so only one static IP address (for the Windows PN proxy `WPY_`<onQhostname>`_`<ESXhostname>) is needed:



Click the **ENROLL** button, then **OKAY**.



The PNs that you selected appear in the Protected Nodes list; however, they are `unverified` as evidenced by the status in the **PROTECTED NODES** page > **Protection Disabled** column.

g. From the ESXi/ESX host, power on the PN proxies.

h. Activate (aka *verify*) the PNs so that onQ can protect them.

Go to **PROTECTION CONFIG** tab, select the PN, then click the **MODIFY** button.

Specify values for the node parameters, then **SAVE**. The act of saving a PN's configuration instructs onQ to activate that PN's configuration.

The onQ Portal groups the PNs by ESXi server. If you do not want these PNs in such a group, clear the default **Group Name** field to remove the PNs from this group and place them in the shared pool.

1.

2. **Create the grub boot menu**. In order to boot the RN successfully, the PN needs to prepare the init ram disk image with the required drivers and legacy grub boot menu.

a. Verify OS version:

```
# cat /etc/redhat-release

Red Hat Enterprise Linux Server release 7.0
(Maipo)
```

b. Make sure `ext2`/`ext3`/`ext4` file systems utilities are installed.

```
# rpm -qa | grep e2fsprogs
e2fsprogs-libs-1.42.9-4.el7.x86_64
e2fsprogs-1.42.9-4.el7.x86_64
```

If not installed, do so now:

```
# yum install e2fsprogs
```

c. Generate the init ram disk with xen drivers and `ext4` file system modules.

Print the kernel release:

```
# uname -r
3.10.0-123.13.2.el7.x86_64
```

Where the `3.10.0-123.13.2.el7.x86_64` is the kernel release by default, change it to match PN's kernel release:

```
# cd /boot
# mkdir -p /boot/grub
# dracut --force --filesystems "ext4 ext3"  \
--add-drivers  "xen:vbd xen:vif" \
initramfs-3.10.0-123.13.2.el7xen.x86_64.img
```

d. Verify the legacy grub boot loader:

```
# vi /boot/grub/grub.conf.xvf5
```

Where the `3.10.0-123.13.2.el7.x86_64` is the kernel release by default, change `vmlinuz` and `initramfs` to match PN's kernel release. The kernel parameters are on a single line. Simply copy and paste from following screen.

Where the `root=UUID=855cd484-3852-4984-b568-ee0408c6b590`, the `855cd`... (UUID) is a temporary placeholder and will be replaced by read "`/`"'s UUID during the RN build. Do not make any changes to this parameter.

For example: The contents of `/boot/grub/grub.conf.xvf5`:

```
default=0
timeout=5
title onQ Red Hat Enterprise Linux (3.10.0-123.13.2.el7.x86_64)
root (hd0,0)
kernel /vmlinuz-3.10.0-123.13.2.el7.x86_64 ro root=UUID=855cd484-3852-
4984-b568-ee0408c6b590 plymouth.enable=0 console=hvc0 loglvl=all
cgroup_disable=memory sync_console console_to_ring earlyprintk=xen
nomodeset net.ifnames=0 biosdevname=0 LANG=en_US.UTF-8
initrd /initramfs-3.10.0-123.13.2el7xen.x86_64.img
```

From `/boot`, validate that `vmlinuz-3.10.0-123.13.2.el7.x86_64` and `initramfs-3.10.0-123.13.2.el7xen.x86_64.img` exist in `/boot` folder as shown in the example below:

```
# ls /boot/vmlinuz-3.10.0-123.13.2.el7.x86_64
/boot/vmlinuz-3.10.0-123.13.2.el7.x86_64
# ls /boot/initramfs-3.10.0-
123.13.2.el7xen.x86_64.img
/boot/initramfs-3.10.0-123.13.2.el7xen.x86_64.img
```

1.

2. **Wait for the RN to build, then perform a self-test**.

   Troubleshooting RN Build or Self-test Problems

   Mistakes with the grub boot menu enforcement can prevent the RN from booting. The following list represents the most common errors.

   - Kernel parameters are not on one single line. Some file editors wrap long parameters.
   - You have a typo in `grub.con` or `grub.conf.xvf5` file name.
   - You have a typo in the kernel file name or the initramfs file name, or these files don't exist.
   - There is a mismatch, on the boot menu, between the kernel versions and the initramfs version. If the kernel's version does not match the contents of initramfs, the RN won't boot.The system could have more than one kernel installed:

     7.0:

     `vmlinuz-3.10.0-123.13.2.el7.x86_64`

should match

```
initramfs-3.10.0-123.13.2.el7xen.x86_64.img
```

6.x:

```
vmlinuz-2.6.32-279.el6.x86_64
```

should match

```
initramfs-2.6.32-279.el6.x86_64.img
```

5.x:

```
vmlinuz-2.6.18-371.el5xen
```

should match

```
initrd-2.6.18-371.el5xen.img.5
```

To find the driver versions packed inside the init ram file system (`initramfs`) of the boot menu: Locate the initramfs and kernel name from the boot menu prepared for the RN (you'll find it under `/boot)`, then use the following command to peek the contents of initramfs. For example:

RHEL 6.x or 7.0:

```
# grep kernel /boot/grub/grub.conf.xvf5
kernel /vmlinuz-3.10.0-123.el7.x86_64  ro
root=UUID=9002ec24-fb30-4d16-8a78-b352a807e82b
plymouth.enable=0 console=hvc0 loglvl=all
cgroup_disable=memory sync_console
console_to_ring earlyprintk=xen nomodeset net.if-
names=0 biosdevname=0 LANG=en_US.UTF-8
# grep initrd /boot/grub/grub.conf.xvf5
initrd /initramfs-3.10.0-123.el7xen.x86_64.img
# lsinitrd /boot/initramsfs-3.10.0-
123.el7xen.x86_64.img|grep modules
rw-r--r--   1 root     root         1446 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.dep
-rw-r--r--   1 root     root         2450 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.dep.bin
-rw-r--r--   1 root     root           52 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.devname
-rw-r--r--   1 root     root        82512 Jun 30
2014 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.order
-rw-r--r--   1 root     root          165 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.softdep
-rw-r--r--   1 root     root        28132 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.symbols
-rw-r--r--   1 root     root        34833 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.symbols.bin
```

RHEL 5.x:

```
# grep kernel /boot/grub/grub.conf.xvf5
kernel /vmlinuz-2.6.18-371.el5xen ro
root=/dev/xvda1 rd_NO_LUKS rd_NO_MD rhgb crashker-
nel=auto rd_NO_LVM
# grep initrd /boot/grub/grub.conf.xvf5
initrd /initrd-2.6.18-371.el5xen.img.5
# zcat /tmp/initrd-2.6.18-371.el5xen.img.5|cpio
-t|grep -E “xen|ext”
16524 blocks
lib/ext3.ko
lib/xennet.ko
lib/xenblk.ko
```

3. **Log on to the PN and open the** following ports on the firewall in order for onQ to communicate with the PN.
   - UDP port `5990`
   - TCP ports `5000` and `5990`

1.

Firewalld:

By default, RHEL 7.0 has introduced a new firewall service, a dynamic firewall daemon known as `firewalld`, instead of iptables service by default; however, the traditional iptables service is supported if installed. For details, see the Red Hat Linux 7 Security Guide. If you choose to disable firewalld, there is no need to configure firewalld firewall rules: simply skip this procedure.

`firewalld` daemon and service and iptables service are using iptables commands to configure the netfilter in the kernel to separate and filter the network traffic. `firewalld` stores it in various XML files in `/usr/lib/firewalld/` and `/etc/firewalld/`.

The firewalld firewall defines how *networks zones* can be used to separate networks into different zones. Based on the level of trust, you can decide to place devices and traffic within a particular network zone. Each mutable network zone can have a different combination of firewall rules.

a. Verify that firewalld is in a running state.

b. Check the service status:

```
[root@RHEL70x64-17-167 services]# systemctl status firewalld.service
firewalld.service - firewalld - dynamic firewall daemon
Loaded: loaded (/usr/lib/systemd/system/fire-walld.service; disabled)
Active: inactive (dead)
```

c. Enable the service, if not already enabled:

```
[root@RHEL70x64-17-167 services]# systemctl enable firewalld
ln -s '/usr/lib/systemd/system/firewalld.service' '/etc/systemd/system/dbusorg.edorapro-ject.FirewallD1.service'
ln -s '/usr/lib/systemd/system/firewalld.service' '/etc/systemd/system
/basic.target.wants/firewalld.service'
```

d. Start the service, if not already running:

```
[root@RHEL70x64-17-167 services]# systemctl start firewalld
```

e. On the PN, find the network interface that is used to communicate with onQ. In this example, that NIC is `ens32`.

```
[root@RHEL70x64-17-167 services]# ifconfig
ens32: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>
mtu 1500
inet 10.20.17.167 netmask 255.255.248.0 broadcast
10.20.23.255
inet6 fe80::250:56ff:fe9d:2121 prefixlen 64 sco-
peid 0x20<link>
ether 00:50:56:9d:21:21 txqueuelen 1000 (Ethernet)
RX packets 7115713 bytes 476287831 (454.2 MiB)
RX errors 0 dropped 149791 overruns 0 frame 0
TX packets 924966 bytes 1305413839 (1.2 GiB)
TX errors 0 dropped 0 overruns 0 carrier 0 colli-
sions 0

lo: flags=73<UP,LOOPBACK,RUNNING> mtu 65536
inet 127.0.0.1 netmask 255.0.0.0
inet6 ::1 prefixlen 128 scopeid 0x10<host>
loop txqueuelen 0 (Local Loopback)
RX packets 10 bytes 980 (980.0 B)
RX errors 0 dropped 0 overruns 0 frame 0
TX packets 10 bytes 980 (980.0 B)
TX errors 0 dropped 0 overruns 0 carrier 0 colli-
sions 0
```

f. For the network interface that you identified above, find the network interface's network zone. In this example, the network zone is `work`.

```
[root@RHEL70x64-17-167 services]# firewall-cmd --get-
zone-of-interface=ens32
work
```

Determine your default zone. In the following example, the default zone is `Public`.

```
[root@RHEL70x64-17-167 services]# firewall-cmd --get-default-zone
public
```

g. Associate the zone(s) with the following firewall rules. The same rules can be applied to many zones as needed. In the following example, dcrm-node service is associated with `work` zone for `ens32`. The `dcrm-node.xml` is located at `/usr/lib/firewalld/services`.

```
[root@RHEL70x64-17-167 services]# firewall-cmd --add-service=dcrm-node --permanent --zone=work
success
```

h. Activate the latest firewall rules:

```
[root@RHEL70x64-17-167 services]# firewall-cmd --reload
success
```

Now the PN can communicate with onQ.

i. Set up the rule for RN on the PN site.

The RN will be equipped with `eth0` interface, so apply the rules to `eth0` interface's zone if different from PN's zone. The PN might not have `eth0` interface; in such a case, the RN's `eth0` will be in the default zone.

Find `eth0` network interface's network zone. In this example, it is *not* set:

```
[root@RHEL70x64-17-167 services]# firewall-cmd --get-zone-of-interface=eth0
no zone
```

Determine your default zone. In this example default zone is `Public`. Since `eth0` has no zone dcm-node is associated with `Public` zone:

```
[root@RHEL70x64-17-167 services]# firewall-cmd --get-default-zone
public
```

j. Associate the zone(s) with the following firewall rules. The same rules can be applied to many zones as needed:

```
[root@RHEL70x64-17-167 services]# firewall-cmd --add-service=dcrm-node --permanent --zone=public
success
```

k. Active the latest firewall rules:

```
[root@RHEL70x64-17-167 services]# firewall-cmd --reload
success
```

Now the RN can communicate with onQ, mainly for self-tests.

I. Confirm the firewall rules. The public zone and work zone have TCP ports (5000/5990) and UDP port 5990 opened in this case.

```
[root@RHEL70x64-17-167 services]# iptables -L -n
Chain IN_public (2 references)
target prot opt source destination
IN_public_log all -- 0.0.0.0/0 0.0.0.0/0
IN_public_deny all -- 0.0.0.0/0 0.0.0.0/0
IN_public_allow all -- 0.0.0.0/0 0.0.0.0/0
Chain IN_public_allow (1 references)
target prot opt source destination
ACCEPT tcp -- 0.0.0.0/0 0.0.0.0/0 tcp dpt:5990 ct-
state NEW
ACCEPT tcp -- 0.0.0.0/0 0.0.0.0/0 tcp dpt:5000 ct-
state NEW
ACCEPT udp -- 0.0.0.0/0 0.0.0.0/0 udp dpt:5990 ct-
state NEW
ACCEPT tcp -- 0.0.0.0/0 0.0.0.0/0 tcp dpt:22 ct-
state NEW
Chain IN_work (0 references)
target prot opt source destination
IN_work_log all -- 0.0.0.0/0 0.0.0.0/0
IN_work_deny all -- 0.0.0.0/0 0.0.0.0/0
IN_work_allow all -- 0.0.0.0/0 0.0.0.0/0
Chain IN_work_allow (1 references)
target prot opt source destination
ACCEPT udp -- 0.0.0.0/0 0.0.0.0/0 udp dpt:631 ct-
state NEW
ACCEPT tcp -- 0.0.0.0/0 0.0.0.0/0 tcp dpt:5990 ct-
state NEW
ACCEPT tcp -- 0.0.0.0/0 0.0.0.0/0 tcp dpt:5000 ct-
state NEW
ACCEPT udp -- 0.0.0.0/0 0.0.0.0/0 udp dpt:5990 ct-
state NEW
ACCEPT tcp -- 0.0.0.0/0 0.0.0.0/0 tcp dpt:22 ct-
state NEW
```

iptables:

a. Verify that udp 5990 and tcp 5000 and 5990 ports are open and above the `REJECT` line:

```
# iptables -L -n | grep -E "5900|5000"
ACCEPT     udp  --  0.0.0.0/0            0.0.0.0/0
udp dpt:5990
ACCEPT     tcp  --  0.0.0.0/0            0.0.0.0/0
state NEW tcp dpt:5000
ACCEPT     tcp  --  0.0.0.0/0            0.0.0.0/0
state NEW tcp dpt:5990
REJECT all -- 0.0.0.0/0 0.0.0.0/0 reject-with
icmp-host-prohibited
```

b. If these ports are not open, open them.

Find the input chain name, which is reported in the line that appears after `INPUT`:

```
# iptables -L -line numbers
Chain INPUT (policy ACCEPT)
num target prot opt source destination
1 RH-Firewall-1-INPUT all -- anywhere anywhere
...
10 REJECT all -- anywhere anywhere reject-with
icmp-host-prohibited
```

Using the line number of the `REJECT` line and the chain name (line `10` and line `1`, respectively, in the step above), insert the onQ ports:

```
# iptables -I RH-Firewall-1-INPUT 10 -m state --
state NEW -p tcp --dport 5990 -j ACCEPT
# iptables -I RH-Firewall-1-INPUT 10 -m state --
state NEW -p tcp --dport 5000 -j ACCEPT
# iptables -I RH-Firewall-1-INPUT 10 -p udp --
dport 5990 -j ACCEPT
```

c. Save and restart iptables.

Afterward, verify that the ports are open and above the `REJECT` line as outlined earlier.

```
# service iptables save
Saving firewall rules to /etc/sysconfig/iptables:
[ OK ]
# service iptables restart
Flushing firewall rules: [ OK ]
Setting chains to policy ACCEPT: filter [ OK ]
Unloading iptables modules: [ OK ]
Applying iptables firewall rules: [ OK ]
Loading additional iptables modules:
ip_conntrack_netbios_n [ OK ]
```

2. **(Oracle database) Install `RMAN` scripts**.

   If the PN has an Oracle database, install the `RMAN` scripts so that onQ can execute a hot backup of your database as outlined in Back up and restore Oracle 11g database on Linux.

(RHEL 6.x) To enroll an agent-less Linux PN:

Use this procedure to enroll an agent-less Linux PN running RHEL 6.x.

1. **Install the PN proxies**:
   a. Enable `ssh` on the ESXi/ESX server.
   b. Reserve two static IP address, one for the Windows PN proxy and one for the Linux PN proxy.
   c. Log on to the HA Appliance's onQ Portal.
   d. Click the **PROTECTION CONFIG** tab > double-plus button (**++**).

      The **Add Protected Nodes via Host** dialog appears.
   e. Provide the credentials and either the hostname or the IP address for the ESXi/ESX host, then **GET LIST**.

      The onQ Portal queries the virtual machines hosted by the ESXi/ESX server, then displays the inventory.

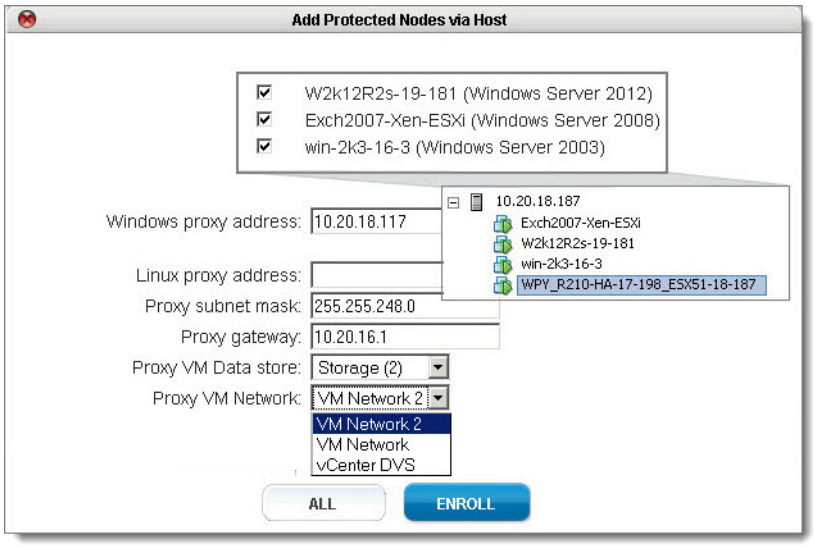
f. In the virtual machine list, select the check boxes (or **ALL** button) for the PNs that you want to enroll.

If you do not see a virtual machine does not appear in the list, it's likely that the virtual machine does not have VMware Tools installed.

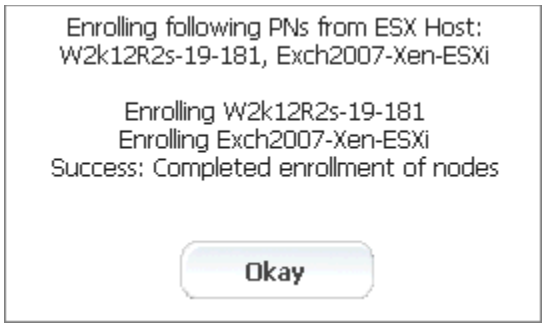
Provide the following information to install the proxies:

- **Windows proxy address**. The static IP address that you reserved for the Windows PN proxy.
- **Linux proxy address**. The static IP address that you reserved for the Linux PN proxy.
- **Proxy subnet mask**. The PN proxy's subnet mask.
- **Proxy gateway**. The PN proxy's gateway.
- **Proxy VM Data store**. The directory in which the vm's files are stored.
- **Proxy VM Network**. Your network configuration can include a vCenter-level Distributed Virtual Switches (DVS) or an ESXi host-level virtual switch. When you enroll an ESXi host, all available networks that are visible to that host are listed and available for selection.

In the following example, three Windows PN are being enrolled, so only one static IP address (for the Windows PN proxy `WPY_`<onQhostname>`_`<ESXhostname>) is needed:



Click the **ENROLL** button, then **OKAY**.



The PNs that you selected appear in the Protected Nodes list; however, they are `unverified` as evidenced by the status in the **PROTECTED NODES** page > **Protection Disabled** column.

g. From the ESXi/ESX host, power on the PN proxies.

h. Activate (aka *verify*) the PNs so that onQ can protect them.

Go to **PROTECTION CONFIG** tab, select the PN, then click the **MODIFY** button.

Specify values for the node parameters, then **SAVE**. The act of saving a PN's configuration instructs onQ to activate that PN's configuration.

The onQ Portal groups the PNs by ESXi server. If you do not want these PNs in such a group, clear the default **Group Name** field to remove the PNs from this group and place them in the shared pool.

1.

2. Copy and modify `/boot/grub/menu.lst`:

```
default=0
timeout=5
splashimage=(hd0,0)/grub/splash.xpm.gz
hiddenmenu
title Red Hat Enterprise Linux (2.6.32-279.el6.x86_64)
    root (hd0,0)
    kernel /vmlinuz-2.6.32-279.el6.x86_64 ro
root=/dev/mapper/VolGroup-lv_root rd_NO_LUKS
LANG=en_US.UTF-8 rd_NO_MD rd_LVM_LV=VolGroup/lv_swap
SYSFONT=latarcyrheb-sun16 crashkernel=auto
rd_LVM_LV=VolGroup/lv_root  KEYBOARDTYPE=pc KEYTABLE=us
rd_NO_DM rhgb quiet
    initrd /initramfs-2.6.32-279.el6.x86_64.img
```

Note: If you want a custom boot menu, create a `/boot/grub/grub.conf.xvf5` file. If `.xvf1` thru `.xvf4` exist, delete them because those files have a higher priority than `.xvf5`.

3. **Wait for the RN to build, then perform a self-test**.

   Troubleshooting RN Build or Self-test Problems

   Mistakes with the grub boot menu enforcement can prevent the RN from booting. The following list represents the most common errors.

   - Kernel parameters are not on one single line. Some file editors wrap long parameters.
   - You have a typo in `grub.con` or `grub.conf.xvf5` file name.
   - You have a typo in the kernel file name or the initramfs file name, or these files don’t exist.
   - There is a mismatch, on the boot menu, between the kernel versions and the initramfs version. If the kernel's version does not match the contents of initramfs, the RN won't boot.The system could have more than one kernel installed:

     7.0:

     `vmlinuz-3.10.0-123.13.2.el7.x86_64`

     should match

```
initramfs-3.10.0-123.13.2.el7xen.x86_64.img
```

6.x:

```
vmlinuz-2.6.32-279.el6.x86_64
```

should match

```
initramfs-2.6.32-279.el6.x86_64.img
```

5.x:

```
vmlinuz-2.6.18-371.el5xen
```

should match

```
initrd-2.6.18-371.el5xen.img.5
```

To find the driver versions packed inside the init ram file system (`initramfs`) of the boot menu: Locate the initramfs and kernel name from the boot menu prepared for the RN (you'll find it under `/boot)`, then use the following command to peek the contents of initramfs. For example:

RHEL 6.x or 7.0:

```
# grep kernel /boot/grub/grub.conf.xvf5
kernel /vmlinuz-3.10.0-123.el7.x86_64  ro
root=UUID=9002ec24-fb30-4d16-8a78-b352a807e82b
plymouth.enable=0 console=hvc0 loglvl=all
cgroup_disable=memory sync_console
console_to_ring earlyprintk=xen nomodeset net.if-
names=0 biosdevname=0 LANG=en_US.UTF-8
# grep initrd /boot/grub/grub.conf.xvf5
initrd /initramfs-3.10.0-123.el7xen.x86_64.img
# lsinitrd /boot/initramsfs-3.10.0-
123.el7xen.x86_64.img|grep modules
rw-r--r--   1 root     root         1446 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.dep
-rw-r--r--   1 root     root         2450 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.dep.bin
-rw-r--r--   1 root     root           52 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.devname
-rw-r--r--   1 root     root        82512 Jun 30
2014 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.order
-rw-r--r--   1 root     root          165 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.softdep
-rw-r--r--   1 root     root        28132 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.symbols
-rw-r--r--   1 root     root        34833 Jan 14
07:06 usr/lib/modules/3.10.0-123.el7.x86_64/mod-
ules.symbols.bin
```

RHEL 5.x:

```
# grep kernel /boot/grub/grub.conf.xvf5
kernel /vmlinuz-2.6.18-371.el5xen ro
root=/dev/xvda1 rd_NO_LUKS rd_NO_MD rhgb crashker-
nel=auto rd_NO_LVM
# grep initrd /boot/grub/grub.conf.xvf5
initrd /initrd-2.6.18-371.el5xen.img.5
# zcat /tmp/initrd-2.6.18-371.el5xen.img.5|cpio
-t|grep -E "xen|ext"
16524 blocks
lib/ext3.ko
lib/xennet.ko
lib/xenblk.ko
```

4. **Log on to the PN and open the** following ports on the firewall in order for onQ to communicate with the PN.
   - UDP port `5990`
   - TCP ports `5000` and `5990`

iptables:

a. Verify that udp 5990 and tcp 5000 and 5990 ports are open and above the `REJECT` line:

```
# iptables -L -n | grep -E "5900|5000"
ACCEPT   udp -- 0.0.0.0/0        0.0.0.0/0
udp dpt:5990
ACCEPT   tcp -- 0.0.0.0/0        0.0.0.0/0
state NEW tcp dpt:5000
ACCEPT   tcp -- 0.0.0.0/0        0.0.0.0/0
state NEW tcp dpt:5990
REJECT all -- 0.0.0.0/0 0.0.0.0/0 reject-with
icmp-host-prohibited
```

b. If these ports are not open, open them.

Find the input chain name, which is reported in the line that appears after `INPUT`:

```
# iptables -L -line numbers
Chain INPUT (policy ACCEPT)
num target prot opt source destination
1 RH-Firewall-1-INPUT all -- anywhere anywhere
...
10 REJECT all -- anywhere anywhere reject-with
icmp-host-prohibited
```

Using the line number of the `REJECT` line and the chain name (line `10` and line `1`, respectively, in the step above), insert the onQ ports:

```
# iptables -I RH-Firewall-1-INPUT 10 -m state --
state NEW -p tcp --dport 5990 -j ACCEPT
# iptables -I RH-Firewall-1-INPUT 10 -m state --
state NEW -p tcp --dport 5000 -j ACCEPT
# iptables -I RH-Firewall-1-INPUT 10 -p udp --
dport 5990 -j ACCEPT
```

c. Save and restart iptables.

Afterward, verify that the ports are open and above the `REJECT` line as outlined earlier.

```
# service iptables save
Saving firewall rules to /etc/sysconfig/iptables:
[ OK ]
# service iptables restart
Flushing firewall rules: [ OK ]
Setting chains to policy ACCEPT: filter [ OK ]
Unloading iptables modules: [ OK ]
Applying iptables firewall rules: [ OK ]
Loading additional iptables modules:
ip_conntrack_netbios_n [ OK ]
```

2. (Oracle database) Install `RMAN` scripts.

If the PN has an Oracle database, install the `RMAN` scripts so that onQ can execute a hot backup of your database as outlined in Back up and restore Oracle 11g database on Linux.

(RHEL 5.x) To enroll an agent-less Linux PN:

Use this procedure to enroll an agent-less Linux PN running RHEL 5.x.

1. **Install the PN proxies**:
   a. Enable `ssh` on the ESXi/ESX server.
   b. Reserve two static IP address, one for the Windows PN proxy and one for the Linux PN proxy.
   c. Log on to the HA Appliance's onQ Portal.
   d. Click the **PROTECTION CONFIG** tab > double-plus button (**++**).

      The **Add Protected Nodes via Host** dialog appears.
   e. Provide the credentials and either the hostname or the IP address for the ESXi/ESX host, then **GET LIST**.

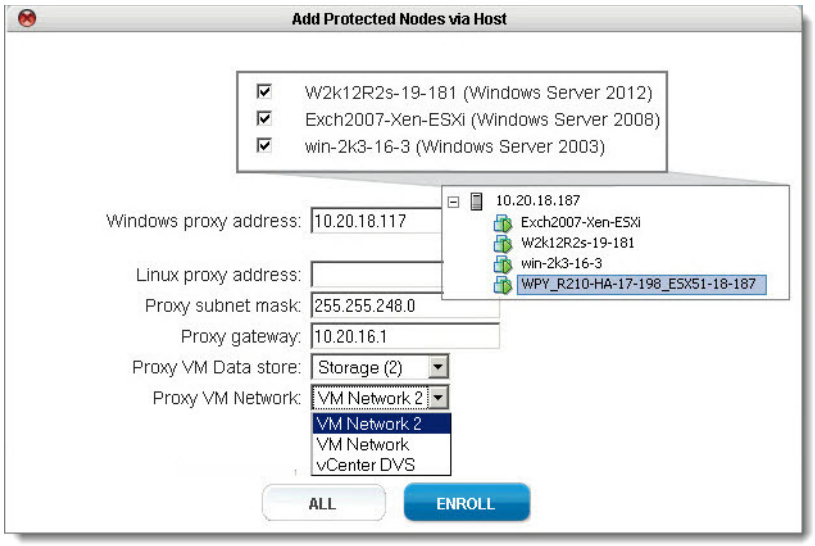
      The onQ Portal queries the virtual machines hosted by the ESXi/ESX server, then displays the inventory.
   f. In the virtual machine list, select the check boxes (or **ALL** button) for the PNs that you want to enroll.

      If you do not see a virtual machine does not appear in the list, it's likely that the virtual machine does not have VMware Tools installed.

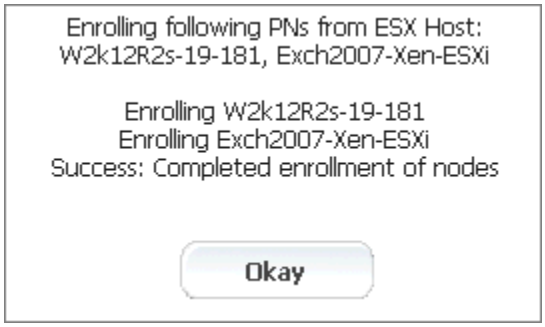
      Provide the following information to install the proxies:
      - **Windows proxy address**. The static IP address that you reserved for the Windows PN proxy.
      - **Linux proxy address**. The static IP address that you reserved for the Linux PN proxy.
      - **Proxy subnet mask**. The PN proxy's subnet mask.
      - **Proxy gateway**. The PN proxy's gateway.
      - **Proxy VM Data store**. The directory in which the vm's files are stored.
      - **Proxy VM Network**. Your network configuration can include a vCenter-level Distributed Virtual Switches (DVS) or an ESXi host-level virtual switch. When you enroll an ESXi host,

all available networks that are visible to that host are listed and available for selection.

In the following example, three Windows PN are being enrolled, so only one static IP address (for the Windows PN proxy `WPY_`<onQhostname>`_`<ESXhostname>) is needed:



Click the **ENROLL** button, then **OKAY**.



The PNs that you selected appear in the Protected Nodes list; however, they are `unverified` as evidenced by the status in the **PROTECTED NODES** page > **Protection Disabled** column.

g. From the ESXi/ESX host, power on the PN proxies.

h. Activate (aka *verify*) the PNs so that onQ can protect them.

Go to **PROTECTION CONFIG** tab, select the PN, then click the **MODIFY** button.

Specify values for the node parameters, then **SAVE**. The act of saving a PN's configuration instructs onQ to activate that PN's configuration.

The onQ Portal groups the PNs by ESXi server. If you do not want these PNs in such a group, clear the default **Group Name** field to remove the PNs from this group and place them in the shared pool.

2. Copy and modify `/boot/grub/menu.lst`:

```
default=0
timeout=5
hiddenmenu
title Red Hat Enterprise Linux Server by Quorum onQ (2.6.18-371.el5xen)
root (hd0,0)
kernel /vmlinuz-2.6.18-371.el5xen ro root=/dev/xvda1 rd_NO_LUKS rd_NO_MD rhgb crashkernel=auto rd_NO_LVM
initrd /initrd-2.6.18-371.el5xen.img.5
```

3. **Log on to the PN and open the** following ports on the firewall in order for onQ to communicate with the PN.
   - UDP port `5990`
   - TCP ports `5000` and `5990`

1.

iptables:

a. Verify that udp 5990 and tcp 5000 and 5990 ports are open and above the `REJECT` line:

```
# iptables -L -n | grep -E "5900|5000"
ACCEPT   udp  --  0.0.0.0/0        0.0.0.0/0
udp dpt:5990
ACCEPT   tcp  --  0.0.0.0/0        0.0.0.0/0
state NEW tcp dpt:5000
ACCEPT   tcp  --  0.0.0.0/0        0.0.0.0/0
state NEW tcp dpt:5990
REJECT all -- 0.0.0.0/0 0.0.0.0/0 reject-with
icmp-host-prohibited
```

b. If these ports are not open, open them.

Find the input chain name, which is reported in the line that appears after `INPUT`:

```
# iptables -L -line numbers
Chain INPUT (policy ACCEPT)
num target prot opt source destination
1 RH-Firewall-1-INPUT all -- anywhere anywhere
...
10 REJECT all -- anywhere anywhere reject-with
icmp-host-prohibited
```

Using the line number of the `REJECT` line and the chain name (line `10` and line `1`, respectively, in the step above), insert the onQ ports:

```
# iptables -I RH-Firewall-1-INPUT 10 -m state --
state NEW -p tcp --dport 5990 -j ACCEPT
# iptables -I RH-Firewall-1-INPUT 10 -m state --
state NEW -p tcp --dport 5000 -j ACCEPT
# iptables -I RH-Firewall-1-INPUT 10 -p udp --
dport 5990 -j ACCEPT
```

c. Save and restart iptables.

Afterward, verify that the ports are open and above the REJECT line as outlined earlier.

```
# service iptables save
Saving firewall rules to /etc/sysconfig/iptables:
[ OK ]
# service iptables restart
Flushing firewall rules: [ OK ]
Setting chains to policy ACCEPT: filter [ OK ]
Unloading iptables modules: [ OK ]
Applying iptables firewall rules: [ OK ]
Loading additional iptables modules:
ip_conntrack_netbios_n [ OK ]
```

2. **(Oracle database) Install RMAN scripts**.

   If the PN has an Oracle database, install the RMAN scripts so that onQ can execute a hot backup of your database as outlined in Back up and restore Oracle 11g database on Linux.

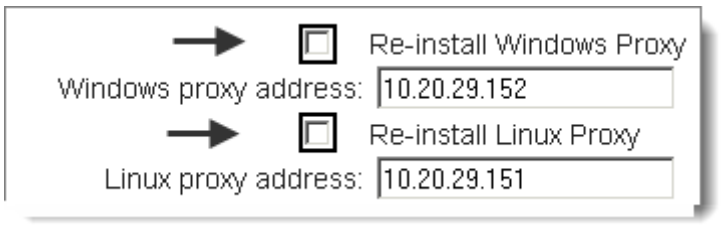
To re-enroll a proxy PN:

Use this procedure to re-enroll (aka reinstall) the Linux PN proxy or the Windows PN proxy or both. Perform this procedure on each ESX/ESXi host enrolled with the onQ.

1. Log on to the HA Appliance's onQ Portal.
2. Click the **PROTECTION CONFIG** tab.
3. Click the double-plus button (**++**).

   The **Add Protected Nodes via Host** dialog appears. The PN check boxes for enrolled PNs are greyed out.
4. Provide the root userid and password to access the ESXi host.
5. In the **Proxy VM Network** field, reattach the virtual network adapter.

6. Select **Re-install Windows proxy** or **Re-install Linux proxy**, depending on the operating system of the proxy you're trying to reinstall, then **ENROLL**.



7. Observe the messages that appear in the pop-up window to verify that the proxy reinstalled successfully.

Related Topics

(Agent-less PNs) Connection and Backup Problems

5.4 (Agent-based PNs) Add protected nodes manually

When you enroll a protected node, onQ automatically (1) installs the onQ Service and (2) adds the PN to the onQ Portal's list of protected nodes thereby registering the PN with the onQ Manager.

However, there are reasons why you might want to *manually* add a protected node to the list:

- You attempted to enroll the protected node, but the installer refused to add the node because there is not enough disk space. Remember, onQ needs disk space for both the protected node's virtual machine (vdisk) and the protected node's data, which is stored in the repository. The amount needed is 1.5 x the used disk space on the protected node. If you are certain that the data will not grow, you can add the node manually.
- You want a protected node to appear in the onQ Portal's list of protected nodes but don't want to enroll the node yet.
- You previously enrolled this PN, but deleted it from the onQ Appliance's configuration.

To add a protected node to the list:

1. Log on to the HA Appliance's onQ Portal.
2. Stop protection.

   You can only add to the list of protected nodes when **Protection** is **OFF**.
3. Click the **PROTECTION CONFIG** tab.
4. Click the plus button (**+**).
5. From the **Add a Protected Node** dialog, type the PN's host name, and specify values for the node parameters, then **SAVE**.

If onQ successfully added the Protected Node, a confirmation message appears:



6. From the **PROTECTION CONFIG** tab, confirm that the newly added protected node appears in the list.
7. If you'd eventually like to protect this node, you must stop protection, then enroll the protected node.
8. Start protection.

6

Enrollment in Windows Cluster Services Environment

- Windows Cluster Services Overview
- (Step 1) Prepare the cluster nodes
- (Step 2) Install onQ Service on cluster
- (Step 3) Configure cluster for the onQ Service
- (Step 4) Configure the onQ Appliance for cluster environment
- (Step 5) Enroll the cluster on the onQ Appliance
- Work with RNs in cluster environment

6.1 Windows Cluster Services Overview

onQ supports the ability to run RNs in a Windows Cluster Services environment. How you enroll and start PNs/RNs is different than in a non-cluster environment. For additional information on cluster support, go to "Windows Cluster Services Support" in onQ Release Notes.

When Exchange or SQL is installed in a cluster environment, it receives its own identity. This identity is not the identity of either of the nodes. This identity is called the Clustered Server (virtual server), and it has its own IP address. This IP address is shared between the nodes. The active node "owns" this virtual IP address until that node fails. The passive node receives heartbeat packets from the active server as long as the active server is up and running.

When the active server goes down/offline, the passive server becomes the active server and the virtual IP address transfers to the new active node. onQ continues to back up the node because it does not see the actual node go offline. In order to support this functionality, the installation of the onQ Service is different than in a non-cluster environment.

Related Topics

Windows Cluster Services Support
Windows Cluster Services Limitations
(Step 1) Prepare the cluster nodes
(Step 2) Install onQ Service on cluster
(Step 3) Configure cluster for the onQ Service
(Step 4) Configure the onQ Appliance for cluster environment
(Step 5) Enroll the cluster on the onQ Appliance
Work with RNs in cluster environment

6.2 (Step 1) Prepare the cluster nodes

To install the cluster modules (*Windows 2012 & 2012R2*):

If your PNs (aka cluster nodes) are running Win2k12 or Win2k12R2, you must perform the following procedure on those PNs prior to protecting the onQ. If you don't, the corresponding RNs will not work.

1. Open a power shell and choose to **Run as Administrator**.
2. Add the required modules.

   Both the Failover Cluster Automation Server module and the Failover Cluster Command Interface module must be installed on each PN (cluster node).

a. Determine which modules are currently installed:

```
> PS C:\Users\administrator> Get-WindowsFeature RSAT-Cluster*
Display Name                                    Name           Install State
------------                                    ----           -------------

[X] Failover Clustering Tools                   RSAT-Clustering
Installed

    [X] Failover Cluster Management Tools       RSAT-Clustering-Mgmt
    Installed

    [X] Failover Cluster Module for Windows... RSAT-Clustering-Powe...
    Installed

    [X] Failover Cluster Automation Server      RSAT-Clustering-Auto...
    Installed

    [ ] Failover Cluster Command Interface      RSAT-Clustering-CmdI...
    Available
```

b. Add the **Failover Cluster Automation Server** module:

```
> Install-WindowsFeature -Name RSAT-Clustering-
AutomationServer
```

c. Add the **Failover Cluster Command Interface** module:

```
> Install-WindowsFeature -Name RSAT-Clustering-
CmdInterface
```

3. Import the modules that you just added:

```
> import-module failoverclusters
```

4. Verify that the required modules are installed:

```
> Get-WindowsFeature RSAT-Cluster*
Display Name                                            Name                    Install State
------------                                            ----                    -------------
[X] Failover Clustering Tools                           RSAT-Clustering
Installed
    [X] Failover Cluster Management Tools               RSAT-Clustering-Mgmt
    Installed
    [X] Failover Cluster Module for Windows...          RSAT-Clustering-Powe...
    Installed
    [X] Failover Cluster Automation Server              RSAT-Clustering-Auto...
    Installed
    [X] Failover Cluster Command Interface              RSAT-Clustering-CmdI...
    Installed
```

Next Step: Go to (Step 2) Install onQ Service on cluster.

To install the iSCSI initiator (*Windows 2003*)

Microsoft iSCSI initiator must be installed on the PN, if not installed already. Installation is required on Windows 2003. The iSCSI initiator is preinstalled on the following Windows versions:

- Windows 2012
- Windows 2008 R2
- Windows 2008

Once installed on the protected node, the initiator will become available in the Recovery Node and will be used for configuration of Virtual SAN connectivity when Recovery Node is started in either Test Mode or in Production Mode.

Go to http://www.quorum.net/file_share/Initiator.zip. The `Initiator.zip` file contains following versions of initiators. Install the appropriate version on the cluster node.

- `Initiator-2.08-build<version>-x64fre.exe`
- `Initiator-2.08-build<version>-x86fre.exe`

Next Step: Go to (Step 2) Install onQ Service on cluster.

6.3 (Step 2) Install onQ Service on cluster

Use this procedure to download and install the onQ Service installer on the Windows cluster server that you want to protect.

To manually install the onQ Service on a cluster server:

Before You Begin: Go to (Step 1) Prepare the cluster nodes.

In this procedure, you will perform a *manual* install of the onQ Service on the cluster server.

1. RDP to the server that you want to protect, then log on to that server as a user with administrative privileges.
2. From that server, launch a browser (see Browser Support), then log on to the HA Appliance's onQ Portal as `varadmin`.
3. Download the onQ Service installer.
   a. In the onQ Portal, go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **DOWNLOADS** page.

Warning: Do not download the `.ez` file, which is only intended for *Protect Me-based* enrollments as outlined in (Agent-based PNs) Restart the onQ Service), not manual installs.

b. Scroll down and select the `QuorumWinBCVSetup<bitVersion>-BCV<build>.msi` file that matches your server's operating system.

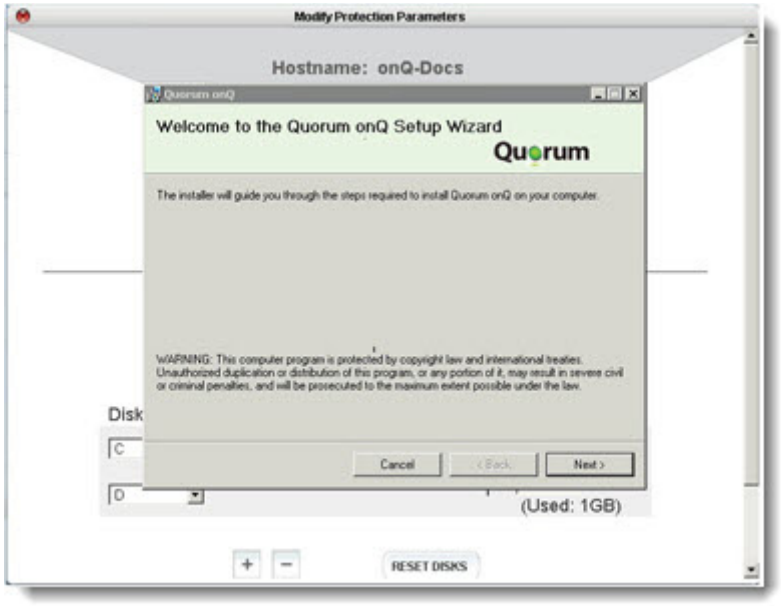
Filename	File ID	Size	Timestamp	Type
QuorumWinBCVSetup64-BCV-3.8-140623-0716.ez	1232	7721011	2014-07-01 11:41:21	Node SW
QuorumWinBCVSetup32-BCV-3.8-140623-0716.msi	1233	6085120	2014-07-01 11:41:21	Node SW
QuorumWinBCVSetup64-BCV-3.8-140623-0716.msi	1234	6250496	2014-07-01 11:41:21	Node SW
DCRMnode-4.5.2BCV-6032.el5.i386.rpm	1235	3681271	2014-07-01 11:41:21	Node SW
QuorumWinBCVSetup32-BCV-3.8-140623-0716.ez	1231	7555635	2014-07-01 11:41:20	Node SW

c. Click the **DOWNLOAD** button, saving the file to your `Downloads` folder. Stop! Do not run this file yet because you might not have the necessary User Account Control levels.

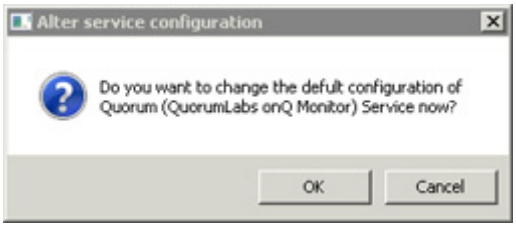
4. Install the onQ Service:

a. Open the `Downloads` folder. On the newly downloaded msi installer, right-click > **Run as Administrator**. If prompted, do not put any restrictions on this account.

b. Follow the on-screen instructions in the **Setup Wizard**.



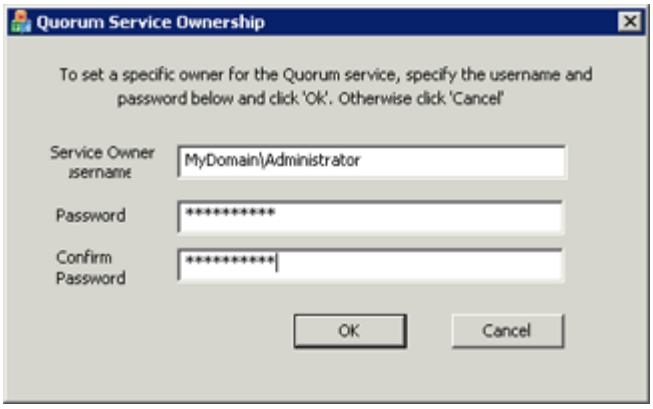
When prompted, choose whether or not you want to configure the service to use a specific account. This account must have adequate privileges to back up all the files that you want to protect.



In the After service configuration dialog, do one of the following:

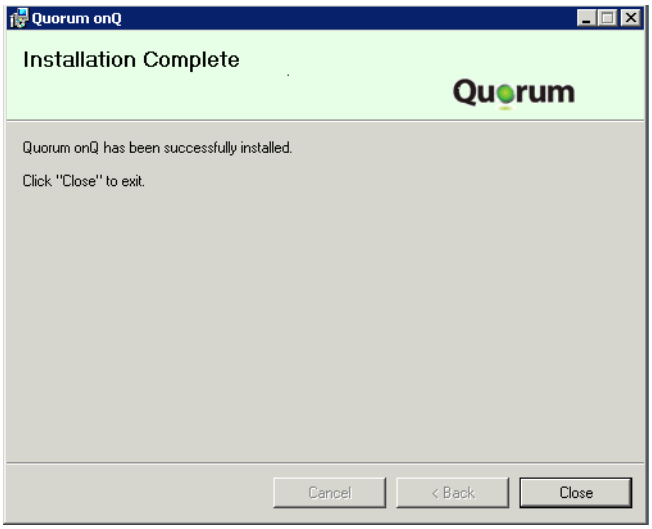
- Click **Cancel** to configure the service to use `Local System` as the owner of the onQ Service.
- Click **OK**, then, in the dialog that launches, specify the account and credentials that you want the service to use, then **OK**. For example, you can configure the service to run as `administrator`. If you want the service to run as the *domain*

administrator specify the domain:



A service running on Windows generally runs as a user account, usually `Local System` by default. However, you can configure the service to run as any user. At the end of the process, onQ Manager presents you with a summary of the PN's proposed configuration.

Congratulations, you've installed the onQ Service.



Next Step: Go to (Step 3) Configure cluster for the onQ Service.

6.4 (Step 3) Configure cluster for the onQ Service

Use this procedure to create a resource and resource group for the onQ Service on the Windows cluster server that you want to protect.

To configure the cluster for onQ Service (*Windows 2012*):

Before You Begin: Go to (Step 2) Install onQ Service on cluster.

Add the “Quorum onQ Service” as a Generic Service Resource:

1. On the active cluster node, launch the Failover Cluster Manager from the **Start** menu and connect to the cluster.
2. Under the Cluster name on the left pane, right-click **Roles** and select **Configure Role**.
3. In the resulting High Availability wizard, do the following:
   a. On the **Before you Begin** page (if shown), click **Next**.
   b. On the **Select Role** page, select **Generic Service** and click **Next**.
   c. Select **Quorum onQ Service** and click **Next**.
   d. On the **Client Access Point** page, do the following:
      - In the **Name** field, type `QuorumOnQSvc`, or type a unique name if onQ is protecting more than one cluster setup and if clusters are using the same Domain Controller to prevent any conflicts.
      - In the **Address** field, type a valid available IPv4 address.
   e. On the **Select Storage** page, click **Next**.
   f. On the **Replicate Registry Settings** page, click **Next**.
   g. On the **Confirm** page, click **Next**.
   h. On the **Summary** page, click **Finish**.

Configure the “Quorum onQ Service” resource:

1. On the active node, launch the Failover Cluster Manager from the **Start** menu and connect to the cluster.
2. Click on the **Roles**.

3. Click **QuorumOnQSvc** in the top middle pane.
4. Switch to the **Resources** tab in the middle pane, at the bottom, and right-click **Quorum onQ Service** under **Roles** and select **Properties**.
5. On the **General** tab, do the following:
   a. Clear the entry in the **Startup Parameters** field.
   b. Enter a space.
   c. Click **Apply**.
6. On the **Policies** tab, do the following:
   a. Select the **if resource fails, attempt to restart...** radio button.
   b. In the **Period for restart(mm:ss)** field, type a value of `2:00`.
   c. In the **Maximum restart in the specified period** field, type `10`.
   d. Select all check boxes on the page.
   e. In the **If all the restart attempts fail, begin restarting again after the specified period(hh:mm)** field, type `01:00`.
   f. In the **Pending timeout(mm:ss)** field, type a value of `03:00`.
7. Click **OK** to exit the wizard.

Restart "QuorumOnQSvc" Resource Group:

1. On the active node, launch the Failover Cluster Manager from the **Start** menu and connect to the cluster.
2. Click **Roles**.
3. Right-click **QuorumOnQSvc** in the middle pane and select **Take offline**.
4. Right-click **QuorumOnQSvc** again and select **Bring online**.

Next Step: Go to (Step 4) Configure the onQ Appliance for cluster environment.

To configure the cluster for onQ Service (*Windows 2008 & 2008 R2*):

Before You Begin: Go to (Step 2) Install onQ Service on cluster.

Add "Quorum onQ Service" as a Generic Service Resource:

1. On the active cluster node, launch the Failover Cluster Manager from the **Start** menu and connect to the cluster.

2. Right-click **Services and Applications** and select **Configure a Service or Application**.
3. In the resulting High Availability wizard, do the following:
   a. On the **Before you Begin** page (if shown), click **Next**.
   b. On the **Select Service or Application** page, select **Generic Service** and click **Next**.
   c. Select **Quorum onQ Service** and click **Next**.
   d. On the **Client Access Point** page, do the following:
      - In the **Name** field, type `QuorumOnQSvc`, or type a unique name if onQ is protecting more than one cluster setup and if clusters are using the same Domain Controller to prevent any conflicts.
      - In the **Address** field, type a valid available IPv4 address.
   e. On the **Select Storage** page, click **Next**.
   f. On the **Replicate Registry Settings** page, click **Next**.
   g. On the **Confirm** page, click **Next**.
   h. On the **Summary** page, click **Finish**.

Configure QuorumOnQSvc Resource group:

1. On the active node, launch the Failover Cluster Manager from **Start** menu and connect to the cluster.
2. Click the **+** to expand **Services and Applications**.
3. Right-click **QuorumOnQSvc** and click **Properties**.
4. On **General** tab, check **Auto Start** (If available).
5. On **Failover** tab, do the following:
   a. In the **Maximum failure in the specified period** field, type a value of `10`.
   b. In the **Period** field, type a value of `1`.
6. Click **OK**.

Configure “Quorum onQ Service” resource:

1. On the active node, launch the Failover Cluster Manager from **Start** menu and connect to the cluster.

2. Click the **+** to expand **Services and Applications**.
3. Click **QuorumOnQSvc**.
4. On the **Summary of QuorumOnQSvc** page (in the middle pane), right-click **Quorum onQ Service**, and click **Properties**.
5. On the **General** tab, do the following:
   a. Clear the entry in the **Startup Parameters** field. If needed enter a couple of spaces.
   b. In the **Please confirm action** window, click **Yes**.
   c. In the **Information** dialog, click **OK**.
   d. Click **Apply**.
6. On the **Policies** tab, do the following:
   a. Select the **if resource fails, attempt to restart...** radio button.
   b. In the **Period for restart(mm:ss)** field, type a value of `2:00`.
   c. In the **Maximum restart in the specified period** field, type `10`.
   d. Select all check boxes on the page.
   e. In the **If all the restart attempts fail, begin restarting again after the specified period(hh:mm)** field, type `01:00`.
   f. In the **Pending timeout(mm:ss)** field, type a value of `03:00`.
7. Click **OK** to exit the wizard.

Restart QuorumOnQSvc Resource group:

1. On the active node, launch the Failover Cluster Manager from the **Start** menu and connect to the cluster.
2. Click the **+** to expand **Services and Applications**.
3. Right-click **QuorumOnQSvc** and select **Take this service or application offline**.
4. Confirm the action.
5. Right-click **QuorumOnQSvc** and select **Bring this service or application online**.

Next Step: Go to (Step 4) Configure the onQ Appliance for cluster

environment.

To configure the cluster for onQ Service (*Windows 2003*):

Before You Begin: Go to (Step 2) Install onQ Service on cluster.

Add QuorumOnQSvc group:

1. On the active node, launch Cluster Administration from **Start** menu and connect to the cluster.
2. Click the **+** to expand **Group**.
3. Right-click **Group** and select New and then Group.
4. In the **New Group** window, do the following:
   a. Type `QuorumOnQSvc`, or type a unique name if onQ is protecting more than one cluster setup and if clusters are using the same Domain Controller to prevent any conflicts.
   b. Enter a description, if desired.
   c. Click **Next**.
5. In the **Preferred Owners** window, do the following:
   a. Select from the **Available Nodes** list and move them to the **Preferred Owners** list, if required. Set the order of preference, if preferred owners are set.
   b. In the **Period** field, type a value of `1`.
6. Click **OK** on the group creation confirmation.

Add “Quorum OnQ Service” as a Generic Service Resource:

1. On the active node, launch Cluster Administration from **Start** menu and connect to the cluster.
2. Right-click **Resources** from under the cluster, select **New**, and then select **Resource**.
3. In the resulting New Resource wizard, do the following:
   a. In the **Name** field, type `Quorum onQ Service`.
   b. In the **Description** field, type the desired description.
   c. In the **Quorum onQ Service** drop-down menu, select **Generic Service**.
   d. In the **Group** drop-down menu, select the **QuorumOnQSvc**.

e. Select the **ON** check box for **Run this resource in a separate resource monitor**.

f. Click **Next**.

4. In the **Possible Owners** window, select desired owners from **Possible Owners** list and click **Next**.
5. In the **Dependencies** window, click **Next**.
6. On the **Generic Service Parameters** page, do the following:
   a. In the **Service Name** field, type `Quorum`.
   b. Leave the **Start Parameters** field empty.
7. In the **Registry Replication** window, click **Finish**.

Configure "Quorum onQ Service" resource:

1. On the active node, launch Cluster Administration from **Start** menu and connect to the cluster.
2. Click the **+** to expand the cluster.
3. Right-click **QuorumOnQSvc** and select **Properties**.
4. On the **QuorumOnQSvc Properties** page, do the following in the **Failover** tab:
   a. In the **Threshold** field, type a value of `10`.
   b. In the **Period** field, type a value of `1`.
   c. Click **OK**.
5. Click **QuorumOnQSvc** on the left pane and select the **Quorum onQ Service** on the right pane.
6. On the **Advanced** tab, do the following:
   a. Select the **Restart** radio button.
   b. Select the **Affect the group** check box.
   c. In the **Threshold** field, type a value of `10`.
   d. In the **Period** field, type a value of `7200`.
   e. In the **Pending timeout** field, type a value `180`.
   f. Click **OK**.

Restart QuorumOnQSvc Resource group:

1. On the active node, launch Cluster Administration from **Start** menu and connect to the cluster.
2. Right-click **QuorumOnQSvc**, and select **Take offline**.
3. Right-click **QuorumOnQSvc**, and select **Bring online**.

Next Step: Go to (Step 4) Configure the onQ Appliance for cluster environment.

6.5 (Step 4) Configure the onQ Appliance for cluster environment

Before You Begin: Go to (Step 3) Configure cluster for the onQ Service.

Use this procedure to configure the onQ Appliance for a cluster environment.

To configure the onQ Appliance for cluster environment:

1. Stop protection.
2. Upgrade onQ to 3.8 SP1 HF1, if not already at that release level.
3. Log on to the hypervisor using Putty or Terminal.
4. Download the SAN template and store it locally on the hypervisor (preferably under `/var/run/sr-mount/xxxxxxxxx-xxxx-xxxx-xxxx` folder).
5. In the console command prompt, do the following:
   a. Extract the SAN template to its original size. When extracted the template should end with `.xva` extension.

```
# gunzip SAN-TEMP.cluster.xva.gz
```

   b. Run the following command, then record the uuid for the item with name-label set to `VHD_SR`:

```
# xe sr-list
```

   c. Using the uuid that you just identified, run the following command:

   The import can take a long time. Nonetheless, the command returns immediately, and so you can monitor the import using

XenServer's log screen; also, when the import disappears from the **ps -ef** command output, the import completed.

```
# xe vm-import filename=SAN-TEMP.cluster.xva sr-uuid=<UUID>
```

To import a template and create the VHD on the hypervisor can take 10-15 minutes.

Next Step: Go to (Step 5) Enroll the cluster on the onQ Appliance.

6.6 (Step 5) Enroll the cluster on the onQ Appliance

Use this procedure to enroll the cluster on the onQ Appliance.

Before You Begin: Go to (Step 4) Configure the onQ Appliance for cluster environment.

To enroll the cluster on onQ:

1. Do the following:
   a. Stop protection.
   b. Go to the **Protection Config** tab.
   c. Click **+** to enroll the cluster.
   d. In the **Hostname** field, enter the cluster name.

      If the cluster name is not known, run the following command on the active cluster node:

```
> cluster /PROPERTIES
```

   e. Set other parameters as needed.
   f. Configure the cluster node and SAN disks that are associated with the cluster. Verify that the **Disk Letter** and corresponding **Size** for each disk is appropriate.
   g. Click **SAVE**.

2. Click on the **Modify** button, then in the protection list, add all volumes and applications that the cluster requires.

   If you are enrolling the PN for the first time using the **Protect Me** button on the onQ Portal, then all volumes (local and iSCSI) on the PN are auto discovered.

3. Click on the **Advanced** button, then in the window, set the radio button for **Enable Cluster Support** to `Yes` to display additional fields for the configuration.

4. From the expanded display for cluster configuration, do the following:

   a. Select the volumes pertaining to cluster services and also any application or other services desired from the list for ‘Cluster Volumes’.

   b. In the **Virtual SAN IP Address** field, type a valid and unique IP address that you want to be used with the virtual SAN. The IP used should be in the same subnet as that of onQ. Also, this IP should not be the same as that of physical SAN currently used with the PN.

   c. Type a gateway address that is associated with the virtual SAN IP address.

   d. Type a subnet mask for the virtual SAN.

5. **SAVE** the protected node configuration.

6. Do the following:

   a. Go to **APPLIANCE CONFIG** tab > **ADVANCED…** button > **HOSTS** tab.

   b. In the **IP Address** field, type the cluster (shared) IP address.

   c. In the **Host(s)/comments** field, type the cluster name.

   d. Click **+** to add the host.

   e. Click **APPLY**.

Next Step: Test your RNs. Go to Work with RNs in cluster environment.

6.7 Work with RNs in cluster environment

After deploying onQ in your cluster environment, use the following procedures to administer your PNs/RNs.

To back up the PN:

1. Start protection.
2. Trigger (scheduled or manual) backups as needed for the configured node.
3. Allow the backup to complete and the RN to be built.
4. If new volumes were added to protection, reinitialize the RN.

   If the PN was added for the first time, this is not required.

To start the RN in test mode (*Windows 2003)*:

Start the Recovery Node for the Domain Controller first in test mode, if the cluster is configured to work with a domain account.

1. Create and start the RN and iSCSI SAN clones. To do so, select the Recovery Node for the cluster node and start it in Test Mode
2. In the RN console for the cluster node, do the following:
   a. Log on with same administrative account that you used on the PN.

      If necessary, to log on to the virtual SAN use the default password for the `root` user account. If you do not know this password, contact Quorum Support.

   b. From a command prompt on the RN, run the following:

```
> cd /d c:\Program Files\Quorum\usr\ONQSUP
> CLUSSERVICEFQ.BAT
```

3. Start the Microsoft iSCSI initiator.
4. Select the **Targets** tab, then log off from any existing or reconnecting targets.
5. Select **Discovery** tab, then do the following:
   a. In the **Target Portals** section, remove any residual entries.

b. Select **Discover Portal** > **ADD**, then type the virtual SAN IP that was provided earlier when you configured the PN in the onQ Portal. Leave the default port set to `3260`.

c. Click **OK**.

6. Select **Targets** tab, then do the following:

a. Select the default Target name or the appropriate one, if more than one exists, then **Logon** button.

b. In the Log on to Target pop-up window, select the **Automatically restore this connection** check box, then **OK**.

c. Select **Persistent Targets** tab, then ensure that the same target that you selected previously on the Targets tab appears in the list.

Note: Click **Finish** to ignore any error while loading a driver for new hardware (IET Controller).

7. Select **OK** on the main window to exit the iSCSI initiator setup.

8. From a command prompt, do the following:

a. Execute the following commands:

```
> cd C:\Program Files\Quorum\usr\ONQSUP
> FIX_CLUSTER.BAT
```

b. Ignore errors at the end of the output and ascertain that the cluster service started.

c. Rejoin the domain, if cluster service fails to start because the RN does not have a domain trust relationship established.

d. Restart the cluster service.

9. Launch Cluster Administrator interface, then do the following:

a. If the Open Connection to Cluster dialog doesn't appear, use the FileOpen Connection dialog.

b. In the Open Connection to Cluster dialog, choose `Open Connection to Cluster` value in the drop-down menu.

c. In the **Cluster or server name** field, type a value '.' (dot), then **OK**.

d. Under **Groups**, ensure that **Cluster Group** and any other **Application Cluster** (SQL Cluster for example) are all on-line, indicating that the cluster is in a usable state.

To start the RN in test mode (*Windows 2008*):

Start the Recovery Node for the Domain Controller in test mode, if the cluster is configured to work with a domain account.

1. Create and start the RN and ISCSI SAN clones. To do so, select the Recovery Node for the cluster node, then start it in Test Mode.
2. In the RN console for the cluster node, do the following:
   a. Log on with same administrative account that you use on the PN.

      If necessary, to log on to the virtual SAN use the password `password` for the `root` user account.

   b. From a command prompt on the RN, run the following:

```
> cd /d c:\Program Files\Quorum\usr\ONQSUP
> CLUSSERVICEFQ.BAT
```

3. Start the Microsoft iSCSI initiator.
4. Select the **Targets** tab, then log off from any existing or reconnecting targets.
5. Select **Discovery** tab, then do the following:
   a. In the **Target Portals** section, remove any residual entries.
   b. Select **Discover Portal** > **Add Portal**, then type the virtual SAN IP that was provided earlier when you configured the PN in the portal. Leave the default port set to `3260`.
   c. Click **OK**.
6. Select **Targets** tab, then do the following:
   a. Select the default Target name or the appropriate one, if more than one exists, then click on the **Logon** button.
   b. In the Log on to Target pop-up window, select the **Automatically restore this connection** check box, then **OK**.
7. Select **Volumes and Devices** tab to list the volumes that reflect the cluster disks. If none appear, choose **Autoconfigure** to force the iSCSI properties to display all cluster volumes configured, then **OK**.

8. From a command prompt, do the following:

   a. Execute the following commands:

```
> cd C:\Program Files\Quorum\usr\ONQSUP
> FIX_CLUSTER.BAT
```

   b. Verify that all disks are marked as `Reserved` using the `diskpart` command. If all volumes expected are not listed and/or are not marked as `Reserved` after few minutes from RN starting, reboot the RN.

9. Rejoin the domain. If the cluster service fails to start because RN does not have a domain trust relationship established, then restart the cluster service.

10. Launch the Failover Cluster Manager interface to perform additional administration, if necessary: **Start** > **Administrative Tools** > **Failover Cluster Manager**.
    - When the interface comes up, the left hand pane lists the clusters available along with the Services and applications, Nodes, Storage, and Networks associated with them.
    - Select the items under **Services and applications**, to list the status of the cluster (server name, disk drives and other resources) as it relates to the specific application. The **Status** of all of these items listed display `on-line` to indicate that the cluster is in a usable state.

To start the RN in test mode (*Windows 2008R2, 2012 & 2012R2*):

Start the Recovery Node for the Domain Controller in test mode, if the cluster is configured to work with a domain account.

1. Create and start the RN and ISCSI SAN clones. To do so, select the Recovery Node for the cluster node, then start it in Test Mode.

2. In the RN console for the cluster node, do the following:

   a. Log on with the administrative account (local or domain if configured for domain) for the PN.

      If necessary, to log on to the virtual SAN, use the password `password` for the `root` user account.

b. From a command prompt on the RN, run the following:

```
> cd /d c:\Program Files\Quorum\usr\ONQSUP
> CLUSSERVICEFQ.BAT
```

3. Start the Microsoft iSCSI initiator.
4. Select the **Targets** tab, then disconnect from any existing targets.
5. Select **Discovery** tab, then do the following:
   a. Remove any residual entries under **Target Portals** and **iSNS servers** sections.
   b. Select **Discover Portal**. Type the virtual SAN IP that you used to configure the PN in the onQ Portal. Leave the default port set to `3260`. Click **OK**.
6. Select **Targets** tab, then do the following:
   a. Select the Target name from under **Discovered Targets**. If the **Status** for the same is `Inactive`, then click on the **Connect** button, bringing the **Status** to `Connected`.
   b. If the Target name doesn't display, type the target IP (same as the virtual SAN IP in the **Discovery** tab), then **Quick Connect** to connect to the virtual SAN, bringing the **Status** to `Connected`.
7. Select **Volumes and Devices** tab to list the volumes that reflect the cluster disks. If none appear, choose **Auto Configure** to force the iSCSI properties to display all cluster volumes configured, then **OK**.
8. From a command prompt, do the following:
   a. Execute the following commands:

```
> cd C:\Program Files\Quorum\usr\ONQSUP
> FIX_CLUSTER.BAT
```

   b. Verify that all disks are marked as `Reserved` using the `diskpart` command. It might take a long time for the volumes to reflect `Reserved` or to appear in the list from `list vol` command. If still not correctly listed after 30 minutes, reboot the RN.
9. If the cluster service fails to start (for example, you cannot see the cluster in the Failover Cluster Manager) or there are issues seeing the domain controller, rejoin the domain.

10. Launch the Failover Cluster Manager interface to perform additional administration, if necessary: **Start** > **Administrative Tools** > **Failover Cluster Manager**.
    - When the interface comes up, the left hand pane lists the clusters available along with Services and applications, Nodes, Storage, and Networks associated with them.
    - Select the items under **Services and applications**, to list the status of the cluster (server name, disk drives and other resources) as it relates to the specific application. The **Status** of all of these items listed display `on-line` to indicate that the cluster is in a usable state.

To start the RN in production mode:

1. Before your start the RN, do the following:
   a. Ensure that physical iSCSI target SAN is turned off and/or the drives configured to work with cluster node are not accessible.

      This is a precautionary step. The Recovery Node registry is updated with the information provided in the onQ Portal for the virtual SAN.

   b. Ensure that all failover nodes in the cluster are turned off.

      The RN you choose must have the latest backup in order for the data to be correct. Only one RN at a time can run in the cluster. When the RN is running in the production cluster, failover is not available.
2. Start the Recovery Node in production mode. Go to “Start recovery nodes on HA Appliance” on page 277 or “Start recovery nodes on DR Appliance or DR Mirror” on page 285.
3. Perform the procedure for your respective platform. The remaining steps are identical to those for test mode.
   - “To start the RN in test mode (Windows 2003):” on page 156
   - “To start the RN in test mode (Windows 2008):” on page 158
   - “To start the RN in test mode (Windows 2008R2, 2012 & 2012R2):” on page 159

7

Updates

- (Start Here) Upgrade Support and Requirements
- Update Appliance software
- (Agent-based Windows PNs) Update node software
- (Agent-based Linux PNs) Update node software
- (Agent-less Windows/Linux PNs) Update PN Proxy software
- Get software packages
- Delete software packages
- (Agent-based Linux PNs) Uninstall node software packages
- Install onQ Appliance license
- View onQ software versions

7.1 (Start Here) Upgrade Support and Requirements

Upgrading *directly* to version 3.8 SP1 HF1 is supported for onQ Appliances and protected nodes running 3.8 SP1.

When upgrading from 3.8 SP1 to 3.8 SP1 HF1, some vulnerability patches are installed and require an additional reboot of onQ. In total, the onQ reboots twice when upgrading from 3.8 SP1 to 3.8 SP1 HF1.

When upgrading, take note of the following:

- When you upgrade your onQ Appliance or AV to 3.8 SP1 HF1 from a version prior to 3.8 SP1, if onQ's or AV's timezone is set to anything other than United States time, the installer resets the timezone such that it shows the incorrect timezone for Backup Status/Archival Status and onQ Status/AV Status. (The timezone in the dialog’s drop-down list is correct.)

To set the Backup Status/Archival Status and onQ Status/AV Status to the correct timezone:

1. Stop protection.
2. Open the **Modify Local onQ Setup**/**Modify Archive Vault Setup** and save the expected timezone: (**APPLIANCE CONFIG** | **ARCHIVE CONFIG** tab > **onQ (LOCAL)** | **ARCHIVE VAULT** page > **MODIFY** button > **SAVE** button).
3. Reboot the onQ and AV.

- After an onQ upgrade to 3.8 SP1 HF1 from a version prior to 3.8 SP1, you must reinstall the Linux PN proxy and the Windows PN proxy on each ESX/ESXi host enrolled with the onQ. This reinstall enables improvements to the PN proxy's handling of power failures. Perform the following steps for each ESX/ESXi host enrolled with the onQ:
  1. Log on to the HA Appliance's onQ Portal.
  2. Click the **PROTECTION CONFIG** tab.
  3. Click the double-plus button (**++**).

     The **Add Protected Nodes via Host** dialog appears.
  4. Provide the root userid and password to access the ESXi host.
  5. Select **Re-install Windows proxy** or **Re-install Linux proxy**, depending on the proxy you're trying to reinstall, then **ENROLL**.
  6. Observe the messages that appear in the pop-up window to verify that the proxy reinstalled successfully.
- After an onQ upgrade, you must refresh your browser using CTRL-F5. If you don't, portal changes will not be activated and some user actions will not work.
- An onQ upgrade to 3.8 SP1 HF1 does not involve any major repository migration or restructuring.
- If the onQ/AV is upgraded from any versions prior to 3.8 SP1, you must manually reboot the onQ Appliance after the onQ upgrade completes.
- Upgrade packages are deployed through onQ Central. onQ and AV will auto-download the necessary onQ packages when the daily cron job runs. This job runs between 10pm and 2am each day. This deployment occurs for all onQ Appliances for which onQ Central is set to allow download of the latest software packages.

Next Step: Now that you've read this topic, you're ready to upgrade. Go to Update Appliance software.

7.2 Update Appliance software

Quorum regularly releases new versions of onQ Appliance software, including service packs and, from time to time, hot fixes.

Upgrade packages are deployed through onQ Central. onQ and AV will auto-download the necessary onQ packages when the daily cron job runs. This job runs between 10pm and 2am each day. This deployment occurs for all onQ Appliances for which onQ Central is set, by Quorum Support, to allow download of the latest software packages.

To *automatically* apply an update:

- You must log on as `VARAdmin` to update the software on your onQ Appliance(s).
- Unless instructed to do so in "(Start Here) Upgrade Support and Requirements" in onQ Administrator's Guide, there is no need to stop protection to apply the updates. If protection is on, onQ quiesces the system, stops protection, applies the new software, then restarts protection. The upgrade process reboots the onQ, if required.
- For onQ Appliances that are multi-tenant, you must upgrade them one tenant at a time and in one session. Unless instructed to do so in (Start Here) Upgrade Support and Requirements, it is not necessary to reboot the onQ Appliance after each tenant upgrade. A reboot after the last upgrade will suffice.
- You must upgrade both HA Appliances, DR Appliances, and onQ Archive Vaults synchronously.

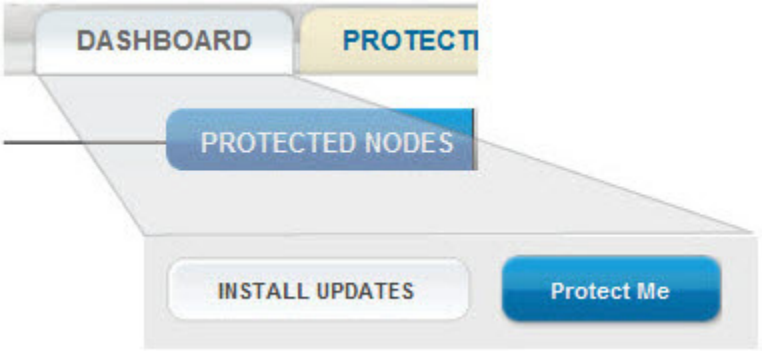
Before you begin, verify that your upgrade path is supported and learn about any pre-upgrade and post-upgrade requirements. Go to "(Start Here) Upgrade Support and Requirements" in onQ Administrator's Guide.

1. Log on to the onQ Portal. You must log on as `varadmin` user.
2. Trigger the download of the onQ package, if it hasn't already downloaded.
   a. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **Software Upgrades** page.

b. Click the **Check for Updates** button.

Mouse-over the Check for Updates button to display the status of the download.

The onQ Portal's **DASHBOARD** displays an **Install Updates** button after the new software downloads from onQ Central. The **onQ STATUS** page > **Issues** pane also displays an `Appliance software upgrade available` message.



3. Apply the new software:
   a. Go to the **DASHBOARD** tab.
   b. Scroll down and click the **INSTALL UPDATES** button.
   c. Click the **UPGRADE** button when prompted to initiate the upgrade.
4. From the **DASHBOARD**, mouse-over the Install Updates icon to display the status of the upgrade.
5. From your browser, press CTRL-F5 to load and activate the onQ Portal changes.

   Immediately after this onQ Appliance upgrade, onQ automatically initiates a PN software (agent-based) and PN proxy software (agent-less) upgrade as outlined in:
   - (Agent-based Windows PNs) Update node software
   - (Agent-based Linux PNs) Update node software
   - (Agent-less Windows/Linux PNs) Update PN Proxy software

To *manually* install an update:

You must log on as `VARAdmin` to update the software on your onQ Appliance(s).

When protection is OFF, the installation process begins immediately. A progress icon appears until the installation completes. Mouse-over the icon to display the log of software packages that are being updated. A reboot of onQ occurs automatically, if needed. Manually installing an update with protection ON is not recommended.

1. Verify that your upgrade path is supported. Go to *(Start Here) Upgrade Support and Requirements*.
1. Log on to the onQ Portal. You must log on as `varadmin` user.
2. Stop protection.
3. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **Software Upgrades** page.
4. Install the files in the following order:
   a. Select onQ upgrade package 1, then **Install**.
   b. Select onQ upgrade package 2, then **Install**. If requested, reboot the onQ Appliance(s).
   c. 'Select 'quark_xxxx.qpkg', then **Install**.
5. Restart Protection.
6. From your browser, press CTRL-F5 to load and activate the onQ Portal changes.

7.3 (Agent-based Windows PNs) Update node software

You do not need to initiate an upgrade of the onQ Service on your agent-based PNs because onQ does this for you automatically and immediately after you initiate an onQ Appliance upgrade and from that point forward if the:

- PN is available.

- onQ is not backing up the PN or processing a recent snapshot from the PN. Otherwise, onQ performs the upgrade after the PN is no longer busy with these tasks.

If you hover-over the **DASHBOARD** tab > **PROTECTED NODES** page > **RN Status** icon, you'll see messages for the upgrade in progress.

Lastly, onQ sends alerts about the upgrade. If you aren't a recipient of such alerts, you can generate the Event Log for a given PN; this log shows the coordination that transpired between onQ and the PNs.

There is no way to disable automatic upgrades of the onQ Service. However, you can use the alternative method outlined in (Alternative) To manually install or reinstall agent-based Windows PN software: if your PNs are in a state that requires a manual upgrade.

Related Topics

(Start Here) Upgrade Support and Requirements
(Agent-based PNs) Protected Node Upgrade Problems
(Agent-based Linux PNs) Update node software
Get software packages

7.4 (Agent-based Linux PNs) Update node software

You do not need to initiate an upgrade of the onQ Service on your agent-based PNs because onQ does this for you automatically and immediately after you initiate an onQ Appliance upgrade and from that point forward if the:

- PN is available.
- onQ is not backing up the PN or processing a recent snapshot from the PN. Otherwise, onQ performs the upgrade after the PN is no longer busy with these tasks.

If you hover-over the **DASHBOARD** tab > **PROTECTED NODES** page > **RN Status** icon, you'll see messages for the upgrade in progress.

Lastly, onQ sends alerts about the upgrade. If you aren't a recipient of such alerts, you can generate the Event Log for a given PN; this log shows the coordination that transpired between onQ and the PNs.

There is no way to disable automatic upgrades of the onQ Service. However, you can use the alternative method outlined in (Alternative) To manually install or reinstall agent-based Windows PN software: if your PNs are in a state that requires a manual upgrade.

To manually update an agent-based Linux PN:

1. Stop protection.

Log on to the Linux PN as `root` or `administrator`.

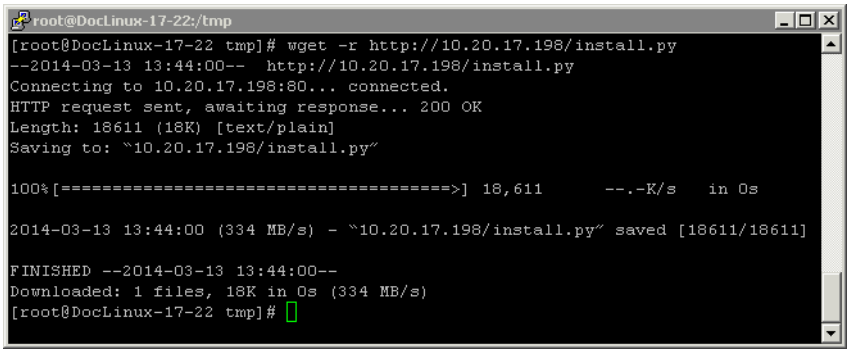
2. Install the agent software:

   a. Launch the installer:

Note: By default, wget preserves the original files. Newly retrieved install.py is saved with extensions `.1`, `.2`, etc. Use the `-r` option (see wget man page) to overwrite the existing file.

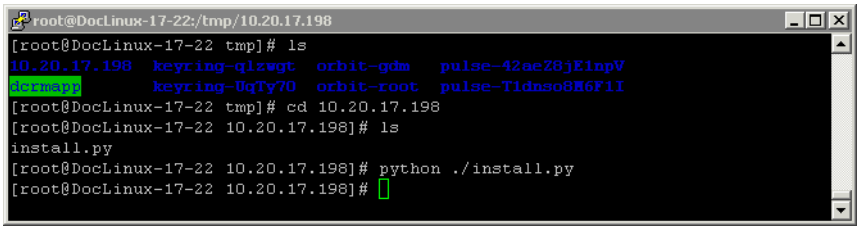
   b. From within a folder (`/tmp`) where you want to save the install script, run the following command:

```
# wget -r http://<onQ-IP-address>/install.py
```



c. Start the installer:

```
# cd <onQ-IP-address>
# python ./install.py
```



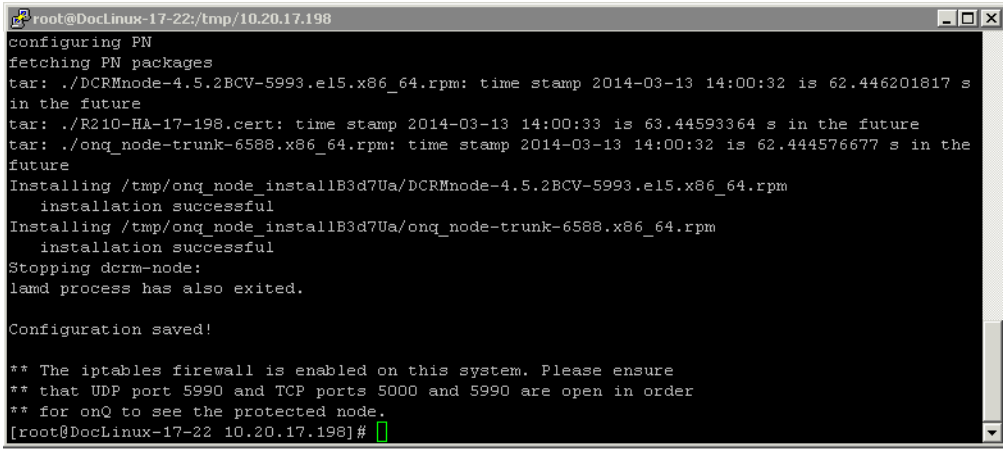
d. Type the credentials for either onQ varadmin or admin user.



e. (Optional) Specify all valid values, excluding file system type and capacity because these values aren't editable, in the node parameter fields provided and modify defaults as needed. (The

install utility lists all available volumes/partitions/mount points and the onQ portal enforces any file system requirements as outlined in Linux RN Limitations.) You can make these node changes at any point after enrollment via the **Protection Config** tab in the onQ Portal.

f. Select **Save** to update/install the client node package on the Linux node.



g. Wait! Do not press Enter. Was the installation successful? Use the error messages below to evaluate the screen output.

- If yes, exit the script.
- If no, press `Enter` to launch the installer UI again. Correct the problem, then repeat Step e through Step g.

`Completed Successfully`
 	
`The iptables firewall is enabled on this system....`
	Message appears in the shell, after the GUI curser exits. In addition, you'll be instructed to open up ports.

`Incorrect/invalid values entered`	Install utility stops if you type incorrect/invalid values. Correct the problem and **Save** again or **Cancel**.
`not authorized`	You either typed the credentials incorrectly or the user account does not have root privileges.

7.5 (Agent-less Windows/Linux PNs) Update PN Proxy software

You do not need to initiate an upgrade of the PN proxies on your ESXi server because onQ does this for you automatically if the ESXi server is available. There is no way to disable automatic upgrades of the PN proxies.

During such an upgrade, vCenter indicates that the PN's proxy, Linux PN proxy (`LPY_`<onQhostname>`_`<ESXhostname>) and Windows PN proxy (`WPY_`<onQhostname>`_`<ESXhostname>), are shutting down and the replacement PN proxies are being deployed.

Also, the **DASHBOARD** tab > **PROTECTED NODES** page > **Connection Status** icon, indicates the upgrade in progress:

1. Before upgrade : solid green carat; white arrow points upward.
2. During upgrade : solid green carat; yellow dash.
3. After upgrade : solid green carat; white arrow points upward.

Also, onQ sends alerts about the upgrade. If you aren't a recipient of such alerts, you can generate the Event Log for a given PN; this log shows the coordination that transpired between the onQ and PN proxies on the ESXi/ESX server.

All PNs on a given ESXi server and having the same operating system (Linux or Windows) will share a common PN proxy. Therefore, when a PN proxy that is common to all PNs is being upgraded, the onQ event log indicates that all the PNs are being upgraded. Also, the **Connection Status** icon provides the same aforementioned status for all the PNs.

7.6 Get software packages

Before you can install onQ Appliance or PN software updates, you might need to do one of the following to get the updates, though it's very unlikely:

- Automatically download packages via onQ Central. Keep in mind that a nightly cron job automatically downloads packages as they are available; however, in the event that a package becomes available before the job is scheduled to run, you'll need to initiate the download to get the latest package.
- Manually upload packages to the onQ Appliance because Quorum Support provides you custom packages to correct an upgrade problem.

After you upload the new software, consider deleting the old version.

To automatically download packages from onQ Central:

Before you can download Appliance software updates, you must enroll with onQ Central, a package distribution system, by installing an Appliance license as outlined in Install onQ Appliance license. onQ queries `https://updates.onqcentral.com`; therefore, ensure that your firewall is configured to allow traffic on outbound port `443`.

After onQ downloads the Appliance software updates, you can manually install them as outlined in Update Appliance software or, in the case of PNs, wait for onQ to install them automatically.

1. Log on to the onQ Portal.
2. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **SOFTWARE UPGRADES** page > **Check for Updates** button.

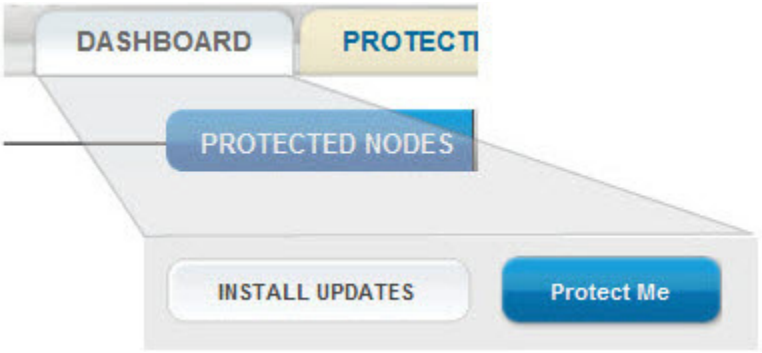
   onQ contacts onQ Central and downloads the latest updates. If a direct upgrade path is not supported for your version, onQ automatically downloads all the prerequisite updates so as to migrate your Appliance and protected nodes to the latest version.
3. Mouse-over the **Check for Updates** button to monitor the status of the download. If there are no new updates, the status indicates `No new updates are avaiable at this time`.

   After onQ downloads the updates, they appear in the list.

Type - indicates `Auto Install` when you download the updates using the method outlined in this procedure.

Installed - indicates when the update was installed on the Appliance. This field is blank until you initiate an install of the update.

Note: The Dashboard displays an **INSTALL UPDATES** button when a package has been downloaded and is available to install. This button remains until you install all available packages:



To manually upload a package to the onQ Appliance:

Although rare, Quorum Support might ask you to upload a special package so as to fix an upgrade problem.

1. From a local computer, download the package from the onQ Download Center.
2. Log on to the onQ Portal.
3. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **Software Upgrades** page > **UPLOAD** button. The Upload SW dialog appears.

   You see a list of a variety of potential software upgrades that have been either delivered with the onQ Appliance or have been transferred to the Appliance or PN subsequent to installation.
4. **Browse** to and select the file that you saved to your local computer.
5. Specify the file type.
   - If the file has a `.qpkg` file format, it's Appliance software. Choose **Appliance SW**.
   - If the file has an `.ez` or `.msi` file format, it's PN software. Choose **Node SW**.
6. Click **Upload**.

Related Topics

Update Appliance software

7.7 Delete software packages

As you download software updates, your list of downloads grows. To keep things tidy and save on disk space, it's a good idea to delete old updates.

There's never a need to revert to an older version of the software; in fact, Quorum discourages this practice.

To delete a node package:

1. Log on to the onQ Portal.
2. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **Downloads** page.
3. Select the file.
4. Click the minus sign (**-**).

To delete an Appliance package:

1. Log on to the onQ Portal.
2. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **Software Upgrades** page.
3. Select the file.
4. Click the minus sign (**-**).

7.8 (Agent-based Linux PNs) Uninstall node software packages

This procedure assumes that you want to uninstall and remove all PN packages.

To uninstall PN packages:

1. Log on to the Linux PN as `root`.
2. From a command line prompt, find and remove the existing node packages.

```
# pkgs=`rpm -qa|grep node`
# echo pkgs=$pkgs
# if [ "$pkgs" != "" ]; then \
rpm -e $pkgs
fi
# pkgs=`rpm -qa|grep node`
# echo pkgs=$pkgs
```

7.9 Install onQ Appliance license

Each onQ Appliance requires a license for each onQ virtual machine (also called, onQ instances) that the onQ Appliance hosts. Licensing is based on the number of onQ instances, enabling support for multi-tenant onQ Appliances, and the number of protected nodes.

Your licenses are managed by and downloaded from onQ Central, a license server and package distribution system. onQ Appliances ships with the perpetual licenses already installed. You'll need to install or reinstall the license on the onQ instance on the onQ Appliance if you want to:

- add a new perpetual license because you've expanded your configuration to include an additional onQ instance on the onQ Appliance.
- enable an onQ instance on the onQ Appliance to manage more protected nodes.
- extend an evaluation license.
- replace an evaluation license with a perpetual license.
- reinstall your perpetual license to fix a license generation problem.

To view the current license:

1. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **LICENSE** page.
2. Observe the following:
   - The serial number appears in the **onQ Instance Serial #** field. You'll need this serial number if you need to request a permanent or replacement license.
   - The number of protected nodes that your license enables you to protect is the **Managed node count**.

To obtain a permanent or replacement license file:

Simply send an e-mail request to license@quorum.net, requesting a replacement license. In that email include the serial number and specify an appropriate return address.

Afterward, Quorum will advise you to do one of the following:

- **Install the license automatically** via onQ Central as outlined in (Recommended) To automatically install a license:.
- **Perform a local installation** as outlined in To manually install a license:. In this case, Quorum will email you the license as an attachment.

(Recommended) To *automatically* install a license:

1. Log on to the onQ Portal as VARadmin.
2. Configure the alerts delivery mechanism. Go to Modify e-mail alert settings.
3. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **LICENSE** page.
4. Click **Install License (QLI)** to check for the latest license on the onQ Central.

5. Wait while the onQ Manager retrieves the latest license from onQ Central, then automatically installs the license on the onQ Appliance.

   When complete, the onQ Manager returns a `License was installed` message. If your onQ Appliance is already running the latest license, the onQ Manager simply reinstalls the same license.

   If you receive either of the following error messages, contact Quorum Support:
   - `Unable to contact QuorumLabs license server...`
   - `Unable to install license...`

To *manually* install a license:

This manual process is not as elegant as allowing the onQ Appliance to do the work for you. Use this manual procedure if onQ Central is not accessible to your onQ Appliance through the internet.

1. Save the license file that Quorum provided in a location that is accessible by your browser.
2. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **LICENSE** page.
3. Click **Install License (Local)** to load and install the license.
4. **Browse...** to the new license file, then **Upload**.

   The new license automatically installs on the onQ Appliance.

Related Topics

License Expiration and Upgrade Alerts

7.10 View onQ software versions

After an upgrade, it's important to view the software status on each onQ Appliance to determine that both Appliances are running the same component versions. If they aren't in sync, upgrade the outdated onQ Appliance. Quorum does not support onQ Appliances running different software versions.

As outlined in How does onQ work?, onQ is comprised of several

components, which are delivered via three distinct packages. onQ Manager is aware of all the versions for each component that a given onQ Appliance is running, but onQ doesn't disclose in the onQ STATUS page the version of the overall installation bundle that the onQ Appliance is running. Only Quorum Support can determine the bundle that the onQ Appliance is running.

To view software status:

1. Log on to the onQ Appliance's onQ Portal.
2. Go to **DASHBOARD** tab > **onQ STATUS** page.
3. Wait a few seconds while the onQ Manager performs a software inventory.
4. Compare the revisions listed in the **Version**, **Platform Status**, **Status** panes against all other onQ Appliances.

7.11 Specify onQ Central Proxy

If all your internet traffic goes through a firewall/proxy server, onQ Central needs to know in order to send updates to the onQ Appliance(s). This *onQ Central proxy* is not to be confused with onQ Proxy.

To specify the onQ Proxy:

1. Log on to the onQ Appliance's onQ Portal.
2. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **Software Upgrades** page.
3. Click the **onQ Central Proxy** button.
4. Specify the networking information for the proxy server, including the IP address, port, and credentials, then **SAVE**.

8

Protection

- Start node protection
- Stop node protection
- Restart node protection
- Modify recovery node creation policy
- Recreate recovery nodes
- Disable protection of nodes
- Enable protection of nodes
- Create on-demand RNs
- Delete on-demand RNs
- Cancel RN updates
- Disable replication globally
- Disable replication for individual nodes
- Enable replication globally
- Enable replication for individual nodes
- Manage hosts
- Resize protected node’s vdisk
- Enable and disable DR Mirroring
- Enroll and Disenroll onQ Archive Vault

8.1 Start node protection

After you enroll protected nodes, you're ready to protect them. The initial backup cycle and recovery node construction will take several hours to complete, but incremental snapshots are usually very fast.

To start protection:

1. Log on to either the HA Appliance's onQ Portal or the DR Appliance's onQ Portal.
2. Select the drop-down menu.
3. Choose **Start Protection**.

   If **Protection** is **ON**, you see the table of nodes you want to protect, with several configurable parameters for each:



8.2 Stop node protection

If you currently have nodes that are being protected and you no longer want to protect them, you must explicitly stop the protection process in order to permanently delete them. The same is true if you want to add a protected node. Also, you must stop protection if you want to change any PN's local or remote configuration or you want to change the hypervisor's configuration.

Stopping the protection process pauses the backup process but preserves existing RNs.

There's no need to stop protection to stop backups for a specific RN. Instead, simply disable protection for that RN.

To stop protection:

1. Log on to either the HA Appliance's onQ Portal or the DR Appliance's onQ Portal.

2. Verify that an RN is not being created, and that a backup is not in progress.

   All icons in **RN Status** column and the **Backup Status** column should be green. It's best practice to wait for these processes to complete before you stop protection.

3. Select the drop-down menu.

4. Choose **Stop Protection**, then **Stop**.

   The **Protection:OFF** icon turns red:



8.3 Restart node protection

From time to time you might need to restart protection to fix a problem with the onQ or a PN update. To restart node protection you simply need to toggle (aka *cycle*) protection: stop protection, then start again.

1. Log on to either the HA Appliance's onQ Portal or the DR Appliance's onQ Portal.

2. Verify that an RN is not being created, and that a backup is not in progress.

   All icons in **RN Status** column and the **Backup Status** column should be green. It's best practice to wait for these processes to complete before you stop protection.

3. Stop protection:

   a. Select the drop-down menu.

   b. Choose **Stop Protection**, then **Stop**.

c. The **Protection:OFF** icon turns red:



4. Start protection:

   a. Select the drop-down menu.

   b. Choose **Start Protection**.

   c. If **Protection** is **ON**, you see the table of nodes you want to protect, with several configurable parameters for each:



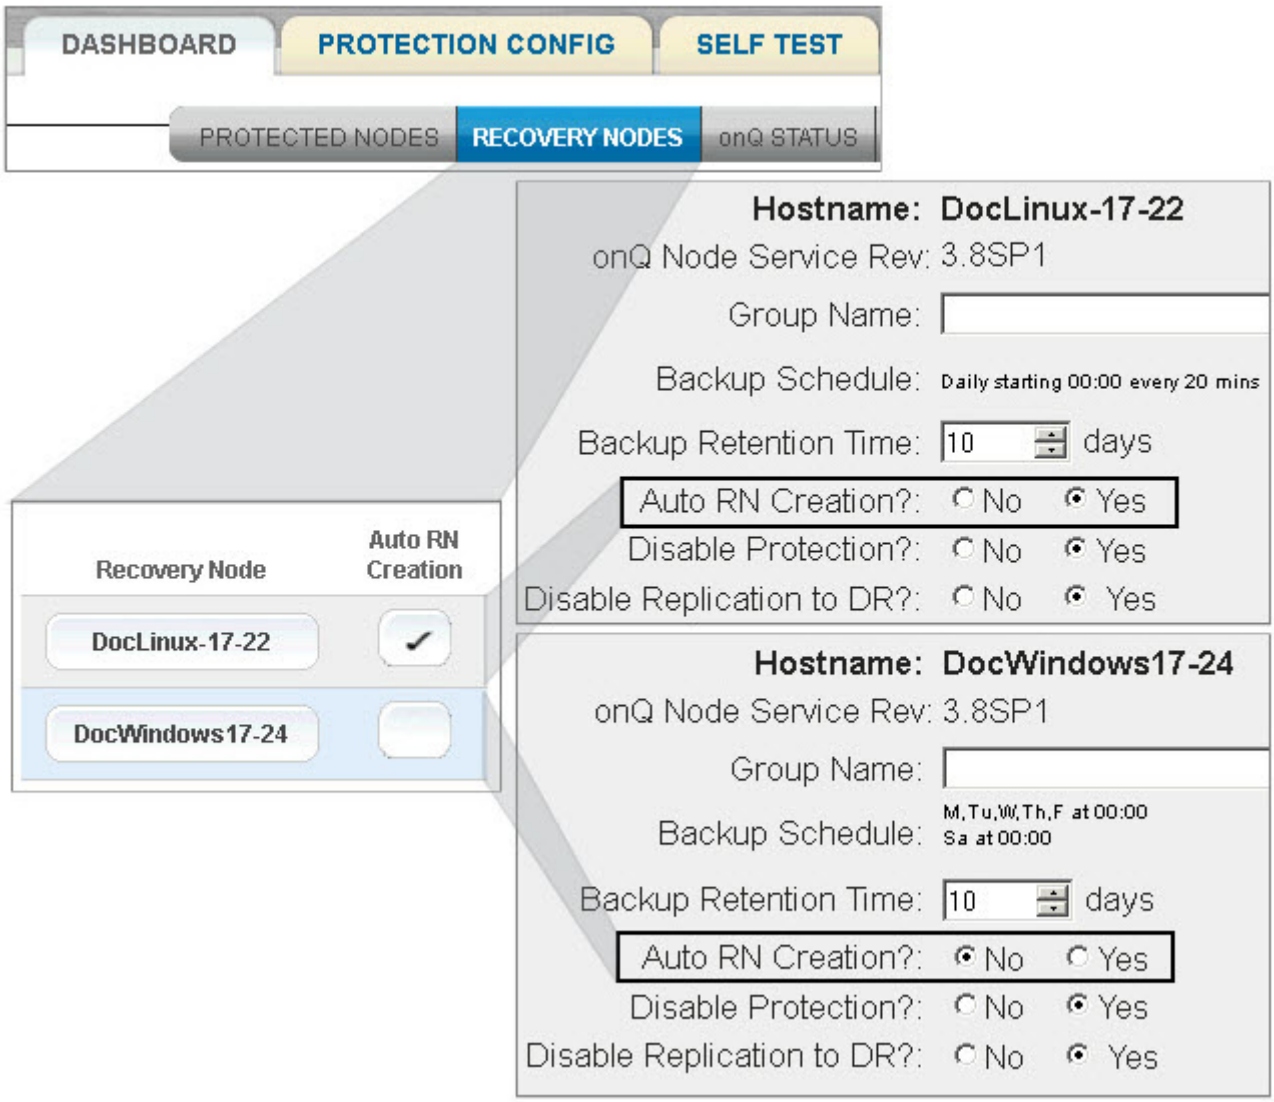
8.4 Modify recovery node creation policy

onQ can create recovery nodes either automatically with each backup cycle or on-demand when requested. This policy represents your potential downtime in the event of a disaster (aka RTO).

When you configure a PN, you must instruct onQ to use one of the following RN build policies:

- **Ready-to-Run** (aka *Instant Disaster Recovery*) - When set to `Auto RN Creation=Yes`, the RN is updated after each backup cycle, and is one-click away from taking the place of the PN. Go to "Start recovery nodes on HA Appliance" on page 277 to see how easy it is to bring an RN online.
- **Build-on-Demand** (aka *Recovery on Demand*) - When set to `Auto RN Creation=No`, the RN is only built when requested through the onQ Portal. This process may take a few hours or more, depending on a few factors. There must be enough disk space on the onQ Appliance to accommodate the RN. The BoD (Build-on-Demand) policy is typically used

to conserve onQ Appliance resources by creating RNs less frequently and only as they are needed, or when you are using the onQ Flex product.



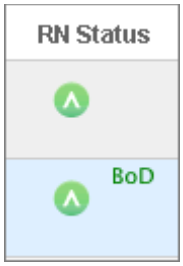
To switch policies:

1. Log on to the HA Appliance's onQ Portal.
2. Go to **DASHBOARD** tab > **RECOVERY NODES** page.
3. Unlock the page.
4. Do one of the following:
   - (All RNs) Click either the **ALL BoD** or **ALL AUTO** button at the

bottom of the page, then click **Disable RN Auto Creation/Enable RN Auto Creation** to switch all RNs to a given policy.

- (Single RN) Click the **Auto RN Creation** button next to the RN, then click **Disable RN Auto Creation/Enable RN Auto Creation** to switch to the other policy.

The acronym `BoD` (Build-on-Demand) appears next to the RN Status icon if `Auto RN Creation` is set to disabled.



Related Topics

Create on-demand RNs

About onQ Flex

8.5 Recreate recovery nodes

When you choose to recreate (aka *reinitialize*/*rebuild*) a recovery node, onQ simply destroys the recovery node and creates a new one from scratch. Recreate a recovery node if:

- You need a new recovery node to replace a corrupt recovery node.
- You need to build a recovery node to test, for example, new Windows Server updates.

Choose any one of the following procedures based on the onQ Appliance that you have available:

To recreate a recovery node (HA Appliance > PROTECTED NODES):

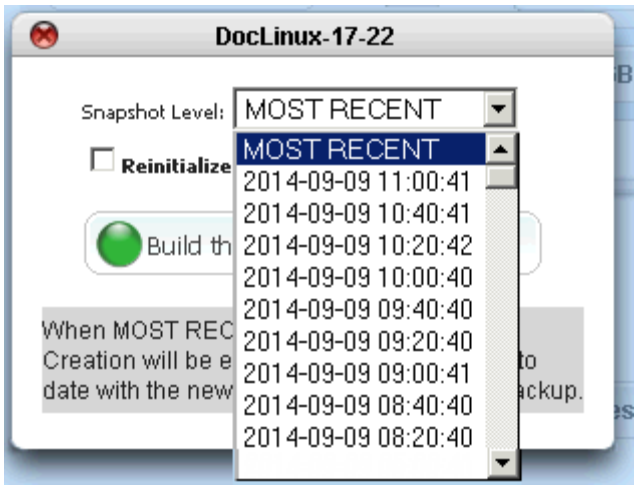
1. Log on to the HA Appliance's onQ Portal.
2. Go to **DASHBOARD** tab > **PROTECTED NODES** page.

3. Unlock the page.
4. Click on the **RN Status** button next to the recovery node, then click **INITIALIZE** and **UPDATE** button. The UPDATE button loads an updated services list (see Edit RN services list).

   The corresponding recovery node is immediately marked as `Not Ready to Run` and the recovery image is recreated from scratch. The **RN status** icon also briefly turns solid red during this process.

To recreate a recovery node (HA Appliance or DR Appliance > RECOVERY NODES):

1. Log on to the HA Appliance's or the DR Appliance's onQ Portal.
2. Go to **DASHBOARD** tab > **RECOVERY NODES** page.
3. Unlock the page.
4. Click on the Protected Node button for the PN that you want to recreate. If there are no snapshots available, this dialog does not appear.



5. Choose a snapshot from the list, then **Build the Recovery Node**. Optionally, you can select the **Reinitialize RN build** check box to rebuild build the RN from scratch instead of updating the RN with incremental data.
   - If you choose `MOST RECENT` snapshot, onQ builds the RN to the most recent snapshot and enables `Auto RN Creation`. If you don't know which one to choose, always choose the latest.
   - If you choose a specific snapshot, onQ builds the RN to the chosen snapshot and disables `Auto RN Creation` to preserve

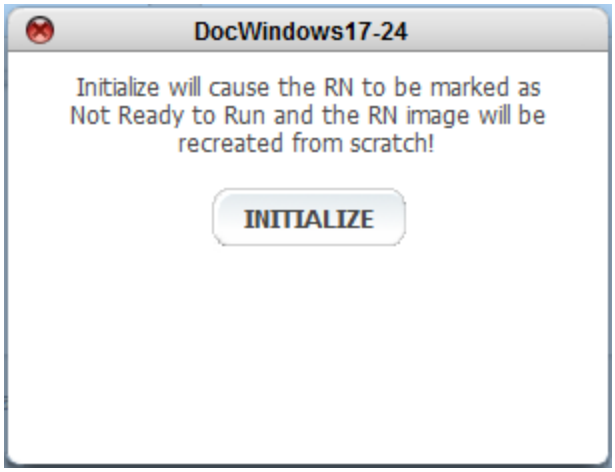
the RN state. Use this option if your latest backup is corrupt. The RN will be built with the current disk configuration for this PN. If the PN's current disk configuration changed significantly (for example, you added a disk/volume) from that of the selected snapshot, the RN won't be functional (might not boot); in this unlikely case, reconfigure the RN to match the PN's current disk configuration.

For more information on the `Auto RN Creation` policy, go to Modify recovery node creation policy.

To recreate a recovery node (DR Appliance):

1. Log on to the DR Appliance's onQ Portal.
2. Go to **DASHBOARD** tab > **DR STATUS** page.
3. Unlock the page.
4. Click on the **RN Ready Status** button next to the recovery node.

   When you roll the cursor over a button in the **RN Ready Status** field, you see a summary of the status of that recovery node as well as the date and time the latest recovery node was generated.
5. In the **Reinitialize RN** information box, click the **INITIALIZE** button and **UPDATE** button. The UPDATE button loads an updated services list (see Edit RN services list).



The onQ Appliance begins the process of rebuilding the image.

Related Topics

Create on-demand RNs

8.6 Disable protection of nodes

If you disable protection, the most recent re-build of the recovery node is retained but backups are suspended.

If you disable protection while a backup is in progress or an RN is being rebuilt, that process continues to completion but further backups are suspended. You can re-enable protection at any time.

To disable protection of a node when *Protection is On*:

1. Log on to the HA Appliance's onQ Portal.
2. Go to **DASHBOARD** tab > **PROTECTED NODES** page.
3. Unlock the page.
4. Do one of the following:
   - (All PNs) Click on the **DISABLE ALL** button at the bottom of the page, then click **DISABLE** to change protection for all PNs.
   - (Group of PNs) For the group, click the **Protection Disabled** ( ) button.
   - (Single PN) Locate the **Protection Disabled** column**.** When you click a button in this column, protection toggles between enabled and disabled for the indicated node.

   A check mark appears in the field if protection is disabled. Otherwise it is blank.

To disable protection of a node when *Protection is Off*:

1. Log on to the HA Appliance's onQ Portal.
2. From the **PROTECTION CONFIG** tab, select the node, then click the **MODIFY** button.
3. Select the **Disable Protection?: Yes** radio button, then **SAVE**.

Related Topics

"Enable protection of nodes" on page 190

8.7 Enable protection of nodes

If you subsequently disabled protection of a PN, you can re-enable at any time. When you re-enable protection, backups resume and subsequent RN rebuilds are based on the retained RN.

To enable protection of a node when *Protection is On*:

1. Log on to the HA Appliance's onQ Portal.
2. Go to **DASHBOARD** tab > **PROTECTED NODES** page.
3. Unlock the page.
4. Do one of the following:
   - (All PNs) Click on the **ENABLE ALL** button at the bottom of the page, then click **ENABLE** to change protection for all PNs.
   - (Single PN) Locate the **Protection Disabled** column**.** When you click a button in this column, protection toggles between enabled and disabled for the indicated node.

   A check mark appears in the field if protection is disabled. Otherwise it is blank.

Related Topics

"Disable protection of nodes" on page 189

8.8 Create on-demand RNs

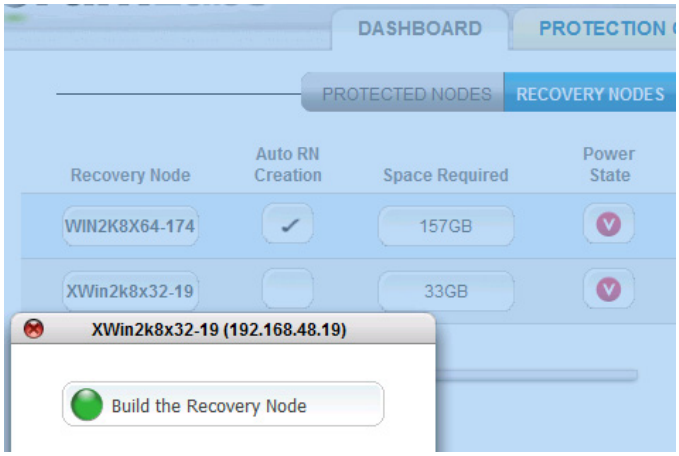
If you specified a Build-on-Demand policy, you must create recovery nodes manually. This process can take 20 minutes to several hours depending on the size of the recovery node. onQ uses the recent backup to create the recovery node. For onQ Flex configurations, there is a cost associated with creating and storing on-demand RNs.

To create a build-on-demand RN:

1. Log on to the HA Appliance's onQ Portal.
2. Go to **DASHBOARD** tab > **RECOVERY NODES** page.
3. Unlock the page.

   The amount of space required appears in the **Space Required** column. The **RN Space Available** is shown just above the Unlock button.
4. Click on the recovery node name, then **Build the Recovery Node**.

   If you don't have enough disk space, onQ states that there's `insufficient disk space`. Go to Delete on-demand RNs.



Related Topics

Modify recovery node creation policy

Delete on-demand RNs

About onQ Flex

8.9 Delete on-demand RNs

Delete Build-on-Demand (BoD) RNs if you need to free up disk space or if you have an onQ Flex configuration and you no longer need the RNs.

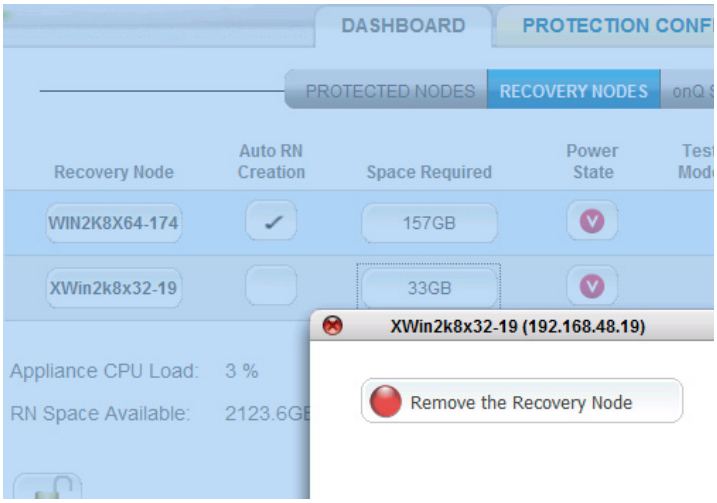
Each RN uses nearly the same amount of disk space on the onQ Appliance as is used on the PN.

Assuming you're using the on-demand policy, if there is not enough disk space on the onQ Appliance for an RN that you want to build, you must remove one or more existing on-demand RNs. (You cannot add disk space at this time without reinstalling.)

Deleting the RN does not affect the corresponding PN.

To remove a recovery node:

1. Log on to the HA Appliance's onQ Portal.
2. Go to **DASHBOARD** tab > **RECOVERY NODES** page.
3. Click on the **Space Required** button for the recovery node that you want to delete.
4. Click **Remove the Recovery Node**.



Related Topics

Create on-demand RNs

Modify recovery node creation policy

8.10 Cancel RN updates

onQ builds Recovery Nodes in the order that they are listed in the onQ Portal. If the first RN in the list has a lot of incremental changes that onQ needs to apply, the other RNs awaiting updates must wait for the first RN update to complete. In this case, you might want to cancel the update for the highest priority RN so as to eliminate the bottleneck.

The RN's next update, which will presumably include even more incremental changes than the update you canceled, will automatically occur as scheduled. If, after the queue clears, you want to immediately apply incremental changes to the RN you can increase its priority by modifying the policy for all but the one RN to BoD as outlined in Modify recovery node creation policy, then restart protection; however, this approach might be overkill.

To cancel a RN update that is in progress:

1. Log on to the HA Appliance's onQ Portal.
2. Go to **DASHBOARD** tab > **PROTECTED NODES** page.
3. Unlock the page.
4. Click on the **RN Status** button next to the recovery node, then click **STOP RN UPDATE** button.



8.11 Disable replication *globally*

Use this procedure to disable backup transfers for all PNs. If you want to disable backup transfers for a given PN, simply disable replication for that PN only.

Quorum does not recommend that you disable backup transfers globally, unless your configuration is an HA-only configuration, meaning you don't have a DR Appliance; if you have an HA-only configuration, disabling all snapshot transfers saves space and improves performance on your HA Appliance.

However, there are situations in which you might need to *temporarily* disable snapshot transfers globally. Before you do so, you must have enough disk space to accumulate backups.

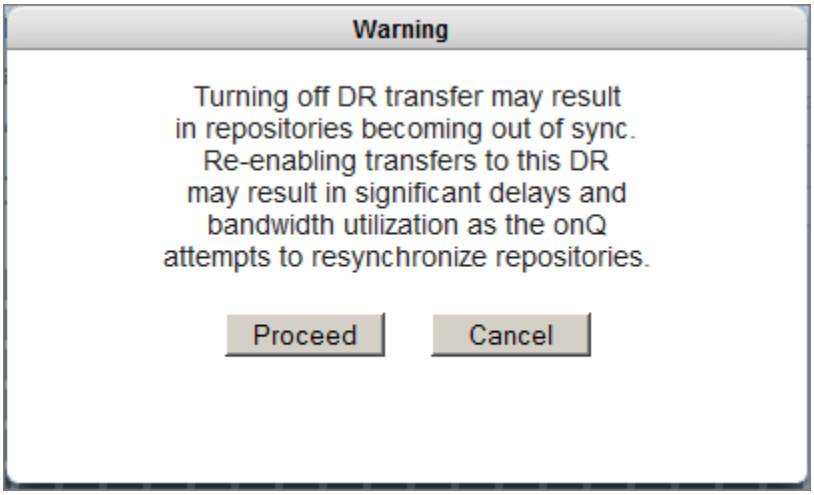
If the HA Appliance cannot communicate with the DR Appliance, it will retain the backups in its queue and transfer them when the DR Appliance comes online.

An alternative to disabling all transfers is stopping protection on the HA Appliance so as to keep the repositories in sync, but the downside is that the HA Appliance will not perform any backups until you turn on protection. This method is not recommended. Besides, when you re-enable transfers, the repositories will resync automatically.

To disable replication *globally*:

1. Log on to the onQ Appliance's onQ Portal.
2. Ensure that your onQ Appliance has enough disk space to accumulate backups. For the related alert, go to Recovery Node and PN Disk Space Alerts.
3. Go to **APPLIANCE CONFIG** tab > **onQ (REMOTE)** page > **MODIFY** button.
4. Do one of the following:
   - (Recommended) Set the bandwidth to zero.
   - Change the DR Appliance's/DR Mirror's IP address to an unused IP address or disconnect the DR Appliance/DR Mirror from the network.
5. Stop protection.

6. In the **Enable DR Transfer?** field, select the **No** radio button, then **Proceed**.



   Disable Transfer: ON icon no longer displays at the top of each UI page on the onQ Appliance.

7. Start protection.

Related Topics

Enable replication globally
Disable replication for individual nodes

8.12 Disable replication for *individual* nodes

Use this procedure to disable transfers for a given PN. If you want to disable transfers for all your PNs, disable replication globally.

You might want to disable replication for a specific PN if it is not important for business continuity, but you do want backups and the ability to archive them.

With protection *enabled* and replication *disabled*, the HA Appliance continues to back up the node, but it will *not* be available on the DR Appliance in the event of a disaster. You can re-enable replication at any time.

To disable replication for a single PN:

1. Log on to the HA Appliance's onQ Portal.

2. From the **PROTECTION CONFIG** tab, select the node, then click the **MODIFY** button.
3. Verify that **Disable Protection?=Yes**. You need protection on to enable PN backups.
4. Select the **Disable Replication?: Yes** radio button, then **SAVE**.

Related Topics

Enable replication for individual nodes
Disable replication globally

8.13 Enable replication *globally*

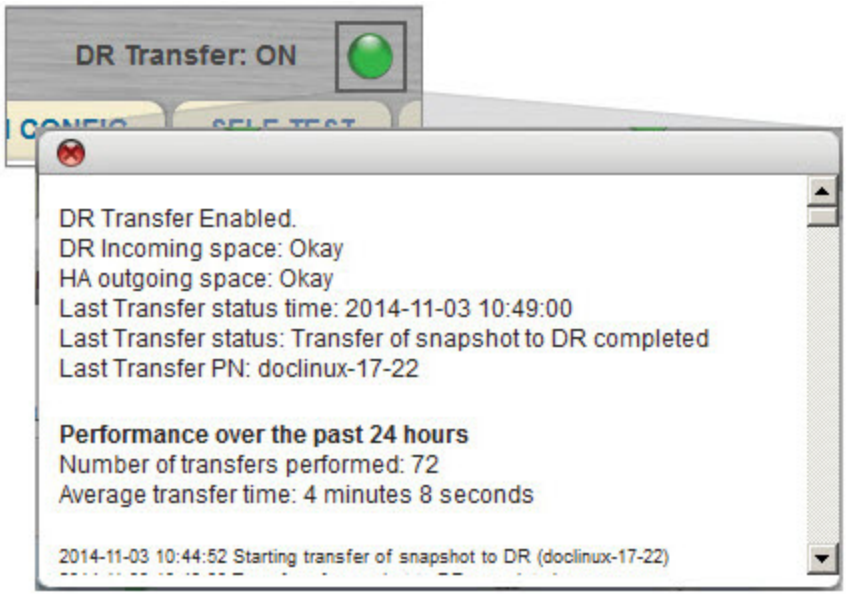
If you subsequently disabled backup transfers globally, you can re-enable at any time. You may experience significant delays and more than usual bandwidth utilization as the onQ resynchronizes the HA Repository and the DR Repository.

To enable replication globally:

1. Log on to the onQ Appliance's onQ Portal.
2. Stop protection.
3. Go to **APPLIANCE CONFIG** tab > **onQ (REMOTE)** page > **MODIFY** button.
4. In the **Enable DR Transfer?** field, select the **Yes** radio button, then **SAVE**.

   Disable Transfer: OFF displays at the top of each UI page on the onQ Appliance.
5. Start protection.

   Disable Transfer: ON displays at the top of each UI page on the onQ Appliance.



Related Topics

Disable replication globally
Disable replication for individual nodes

8.14 Enable replication for *individual* nodes

If you subsequently disabled replication for a specific PN, you can re-enable at any time.

To re-enable replication for a single PN:

1. Log on to the HA Appliance's onQ Portal.
2. From the **PROTECTION CONFIG** tab, select the node, then click the **MODIFY** button.
3. Select the **Disable Replication?: No** radio button, then **SAVE**.
4. Initiate two immediate backups, and wait for two successful backup transfers.

   onQ automatically restores the full snapshot of that PN and transfers that snapshot and the delta snapshot for the next backup to the DR. However, onQ does so based on the PN's preconfigured backup schedule. Quorum recommends that you leapfrog the onQ by executing the backup immediately so as to put the DR into a healthy state as soon as possible.

Related Topics

Disable replication for individual nodes
Enable replication globally

8.15 Manage hosts

An onQ Appliance maintains a table that enables the onQ Appliance to resolve host names to their IP address. Quorum recommends that you include all protected nodes as well as any external nodes. Define these nodes on both the HA Appliance and the DR Appliance. This best practice ensures hostname resolution in the event that your DNS server fails or if you

have a site failure.

To add a host:

1. Log on to the Appliance's onQ Portal.
2. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **HOSTS** page.

   You see a list showing IP address and names of hosts that cannot be resolved by DNS.
3. Click the plus button (**+**).
4. Type the IP Address and hostname in the fields provided, then **SAVE** (or **RESET** to discard your changes).

To delete a host:

If you've removed a node from protection because it's no longer in your production environment, you can delete it from the list of hosts, to keep things tidy. However, if you intend to add the node later, there's really no harm in leaving the host in the list.

1. Log on to the Appliance's onQ Portal.
2. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **HOSTS** page.
3. Select the host form the list, then click the minus button (**-**), then **YES** to confirm.

To edit an existing host:

1. Log on to the Appliance's onQ Portal.
2. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **HOSTS** page.
3. Select the host from the list, then **MODIFY**.
4. Change the IP address or hostname in the fields provided, then **SAVE**.

8.16 Resize protected node's vdisk

As outlined in "(Agent-based PNs) Add protected nodes manually" on page 136, a protected node's virtual disk (vdisk) requires a specific amount of disk space.

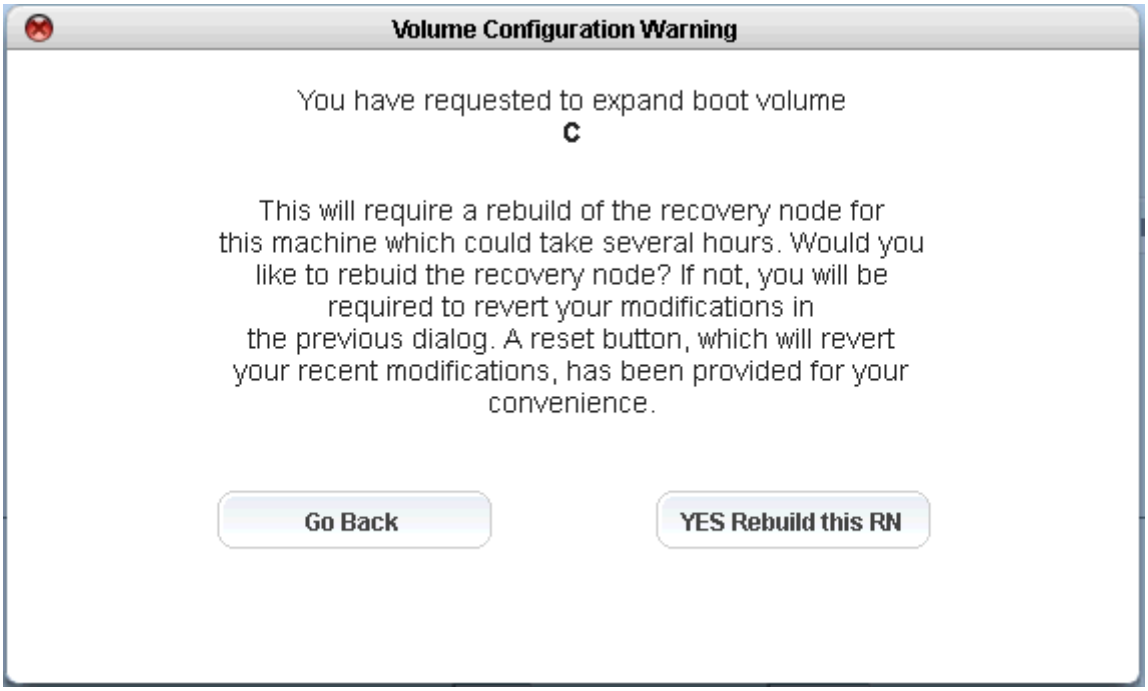
onQ alerts you (see "Recovery Node and PN Disk Space Alerts" on page 417) when your protected node's vdisk capacity used exceeds 90%. If you wait until a vdisk fills up, onQ cannot protect that protected node.

- If you decrease the size of the vdisk for any volume, you must rebuild the entire RN.
- If you increase or decrease the size of `C:\`, you must rebuild the entire RN.
- If you increase the size of a vdisk for any drive/volume other than `C:\`, you do *not* need to rebuild the RN because onQ simply adds an additional vdisk for the delta. However, there is one exception. An RN is limited to a total of 15 vdisks. If you exceed this limit, the onQ Portal prompts you to rebuild the RN so that onQ can consolidate the vdisks and provide you the vdisk size that you requested.

To resize a vdisk:

1. Log on to the HA Appliance's onQ Portal.
2. Go to **PROTECTION CONFIG** tab, select the protected node, then click **MODIFY**.
3. Specify a new size for the vdisk, then **SAVE**.

If prompted, click **YES Rebuild this RN**.



Related Topics

Recovery Node and PN Disk Space Alerts
Monitor disk space and memory usage

8.17 Enable and disable DR Mirroring

A typical onQ configuration consists of an HA Appliance and a DR Appliance, providing high availability. However, if you prefer additional protection in the event that your HA Appliance and your DR Appliance are unavailable, you can expand your existing configuration to include a DR Mirror.

DR mirroring enables you to replicate the data in the DR repository to another onQ Appliance, a *DR Mirror*. A DR Mirror is identical to the DR Appliance. The data moves from the HA Appliance to the DR Appliance, as usual. The only difference is that, with DR mirroring, the DR Appliance transfers/replicates the data on to the DR Mirror.

In a typical DR mirroring configuration, there are 3 different appliances: HA Appliance, DR Appliance, *and* DR Mirror. However, you can "daisy-chain" as many DR Mirrors as you would like; in this case, the DR Appliance is chained to the first DR Mirror, and the first DR Mirror is chained to the second DR Mirror (and so on) by way of trust relationships.

After an initial configuration, do not disable DR mirroring as outlined below.

To enable DR mirroring:

1. Log on to the DR Appliance or DR Mirror's onQ Portal.
2. Go to **APPLIANCE CONFIG** tab > **onQ (REMOTE)** page.
3. Click **Modify**. The **Modfiy Remote onQ Setup** page appears.
4. Select the **Enable DR Mirroring?: Yes** radio button.
5. In the drop-down text boxes provided, specify the networking parameters for the DR Mirror.

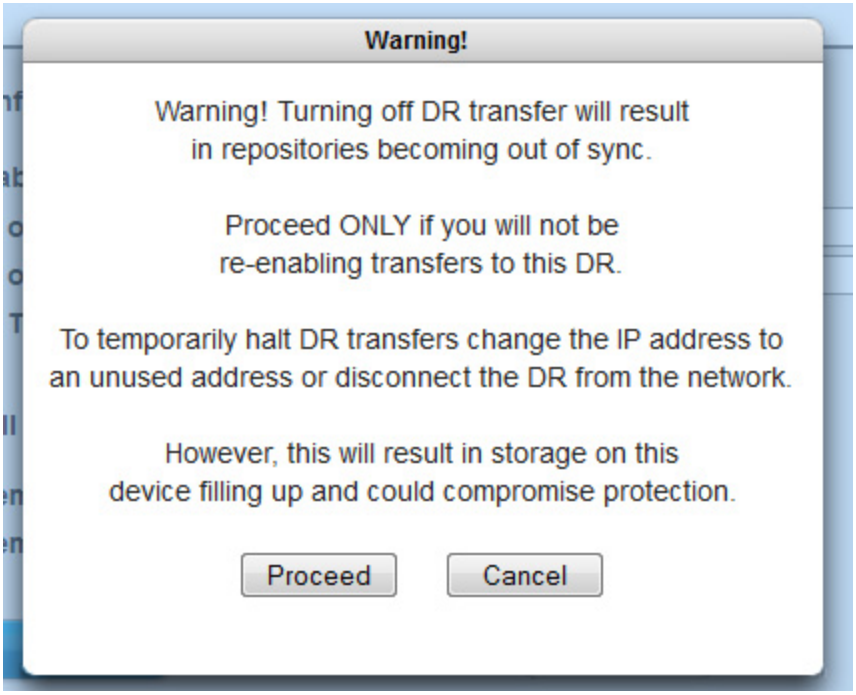
   Mirror onQ is the DR Mirror that you want to mirror. A DR Mirror can mirror a DR Appliance or, in the case of multiple DR Mirrors, the next DR Mirror in the chain.
6. Click **SAVE**.

To disable DR mirroring:

Before you disable DR mirroring, ensure that your onQ Appliance has enough disk space to accumulate backups.

If you need to temporarily stop transfers, change the DR Mirror's IP address to an unused address or disconnect the DR Mirror from the network, as the

warning below advises:.



If the DR Appliance cannot communicate with the DR Mirror, it will retain the backups in its queue and transfer them when the DR Mirror comes online.

Stopping protection on the HA Appliance will also keep the repositories in sync, but the downside is that the HA Appliance will not perform any backups until you turn on protection. This method is not recommended.

Related Topics

Step 2: Configure the onQ Appliances
Set up trust relationships

8.18 Enroll and Disenroll onQ Archive Vault

When you disenroll an onQ Archive Vault on the onQ Appliance, the AV disables all the jobs that correspond to that onQ Appliance.

To learn about enrollment, go to Monitor onQ Archive Vault enrollment or refer to the AV online help.

To disenroll onQ Archive Vault:

1. Log on to the onQ Appliance's onQ Portal.
2. Go to **APPLIANCE CONFIG** tab > **ARCHIVE VAULT** page.
3. Click on the **Disenroll** button.

   The **Disenroll** button now reads **Info**, and the **Archive Enrollment Status** is `Not Enrolled`. In addition, the Archive Vault Portal indicates that the onQ Appliance is disabled.

Related Topics

Monitor onQ Archive Vault enrollment

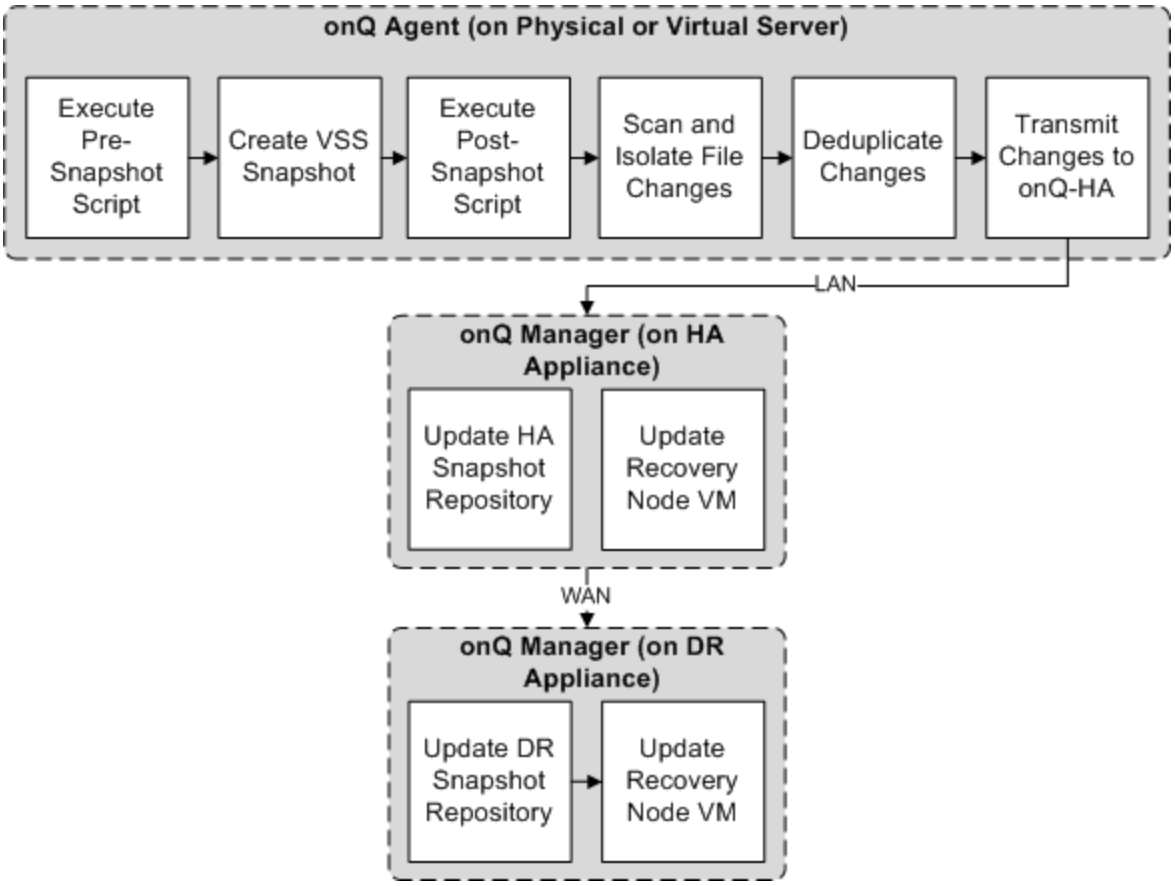
9

Backup and Restore

- Backup and Recovery Workflow
- Edit RN services list
- Edit backup exclude list
- Edit backup include list
- Initiate immediate backups
- Export or import Repository
- Schedule backups
- Change Backup Retention Policy
- Delete snapshots
- Stop in-progress backups
- Retrieve objects from Repository
- Perform WSR Restore
- Perform File-level Restore
- Adjust bandwidth throttling
- Limit resource utilization on PNs
- Run custom backup scripts
- (Agent-based PNs) Methods of Performing Incremental Backups
- Perform full rescan on PN
- Enable onQ to back up shared volumes

9.1 Backup and Recovery Workflow

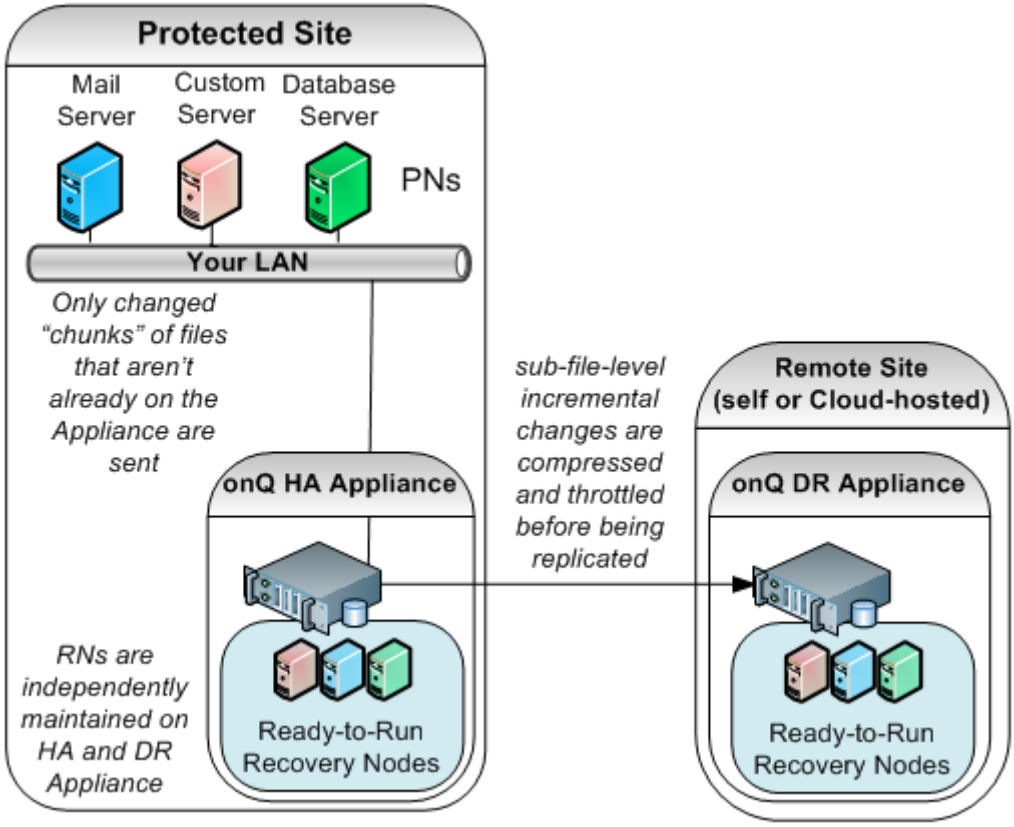
The repository is the deduplicated archive of all snapshots of all of your Protected Nodes. From the repository, onQ can reproduce a complete system image from any snapshot.



onQ provides HA and DR protection in the following way:

- At scheduled intervals, the onQ Manager takes crash-consistent, open-file incremental snapshots of your Protected Nodes.
- The snapshot is deduplicated, so only the changed files (or changed parts of large files) are sent to the HA Appliance.
- The HA Appliance adds the changes to the repository and then compresses and sends the changes to the DR site.
- Ready-to-run recovery nodes are created on the HA Appliance.

- Ready-to-run recovery nodes are created on the DR Appliance.
- The onQ Manager monitors the health of your PNs. If the onQ Manager detects a PN or backup failure, the onQ Manager issues an email alert. You can then determine if the condition of the PN requires that you start the HA or DR recovery nodes using the onQ Portal.



Related Topics

Initialize the DR Repository
Initiate immediate backups
Resize protected node's vdisk

9.2 Edit RN services list

A Recovery Node (RN) services list is a list of services that can run on a given PN. Some services installed on a PN will cause problems if they attempt to run on the RN. In this case, you can configure the startup settings for these services.

The onQ Portal provides a default RN services list that includes all known services that are problematic. For this reason, the start option is set to `Disabled`. Alternatively, the onQ Portal enables you to retrieve a comprehensive list of services from which you can build your own custom list.

- If a custom list exists for an RN, the onQ Portal uses that list instead of the defaults list.
- If you later delete the PN, the onQ Portal deletes the custom list for that PN.

To specify the start option for a service:

This procedure assumes that you're modifying the RN services list for an existing protected node. However, you can perform this procedure from either the **Add Protected Node** page or the **Modify Protected Node** page.

1. Log on to the HA Appliance's onQ Portal.
2. Click the **PROTECTION CONFIG** tab.
3. Select the node, then **MODIFY** button > **ADVANCED** button.
4. In the **RN Services** field, click **Custom** or **Default**.

When you click **Custom**, the onQ Portal overwrites the values from the custom list with the values from the default list so that the known services that are problematic always default to `Disabled`.



5. Do any of the following to create or modify a custom list:

Note: If you delete all entries in the custom list, the onQ Portal assumes you no longer want a custom list and applies the default list.

- Select an existing service from the table and **Modify** its start option.
- Click **Extract** to retrieve a list of services and settings that recently appeared in your PN's registry. Prune and or selectively modify this list to meet your needs. This option prevents you from unintentionally misspelling a service name. If no snapshot exists for the PN, which only occurs if you added a PN and the backup is in progress, then the onQ Portal displays an error.
- In the text box, type the name of the service and specify its start option, then click the plus (**+**) button.

6. Do one of the following:

Note: The onQ Portal validates the service to the extent that it checks for duplicates and blank entries.

- Click **Save** - onQ Portal saves your exclude list as shown in the

list, then returns you to the previous page. The **Backup Exclude List:** button is labeled **Custom**.

- Click **Default** - onQ Portal deletes the custom list and replaces the custom list with the default list (also known as *factory default*), then returns you to the previous page. The **Backup Exclude List:** button is labeled **Default.**
- Click **Revert** - onQ Portal discards your changes, then returns you to the previous page. The **Backup Exclude List:** button does not change.

9.3 Edit backup exclude list

A backup exclude list is a list of folders and files that need not be backed up. The purpose of the list is to improve the efficiency of the backup process by minimizing the quantity of data to be processed. For each protected node there is a backup exclude list.

By default, the onQ Appliance excludes temporary files. You must take great care in selecting files and folders to exclude from backup. Excluding some key files might prevent your recovery nodes from working properly.

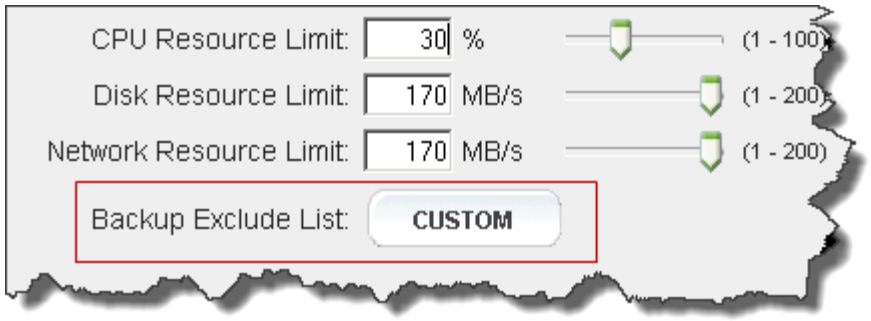
For an exclude list, Quorum recognizes file- and folder-names expressed as IEEE POSIX Basic Regular Expressions (BRE). You will find a synopsis of BRE syntax at ATP://en.wikipedia.org/wiki/Regular_expression.

To add a file or folder to the exclude list:

This procedure assumes that you're modifying the exclude list for an existing protected node. However, you can perform this procedure from either the **Add Protected Node** page or the **Modify Protected Node** page. Exclude lists are specific to the protected node's operating system.

1. Log on to the HA Appliance's onQ Portal.
2. Click the **PROTECTION CONFIG** tab.
3. Select the node, then **MODIFY** > **ADVANCED**.

4. In the **Backup Exclude List** field, click **Custom** or **Default.**



Note: On a DR Appliance, the Backup Exclude List is a read-only field.

5. Specify a new exclusion.
   a. In the text box, type the path to the file or folder. Your path cannot contain any of the following characters:

      < (less than)

      > (greater than)

      : (colon)

      " (double quote)

      / (forward slash)

      \ (backslash)

      | (vertical bar or pipe)

      ? (question mark)

      * (asterisk)

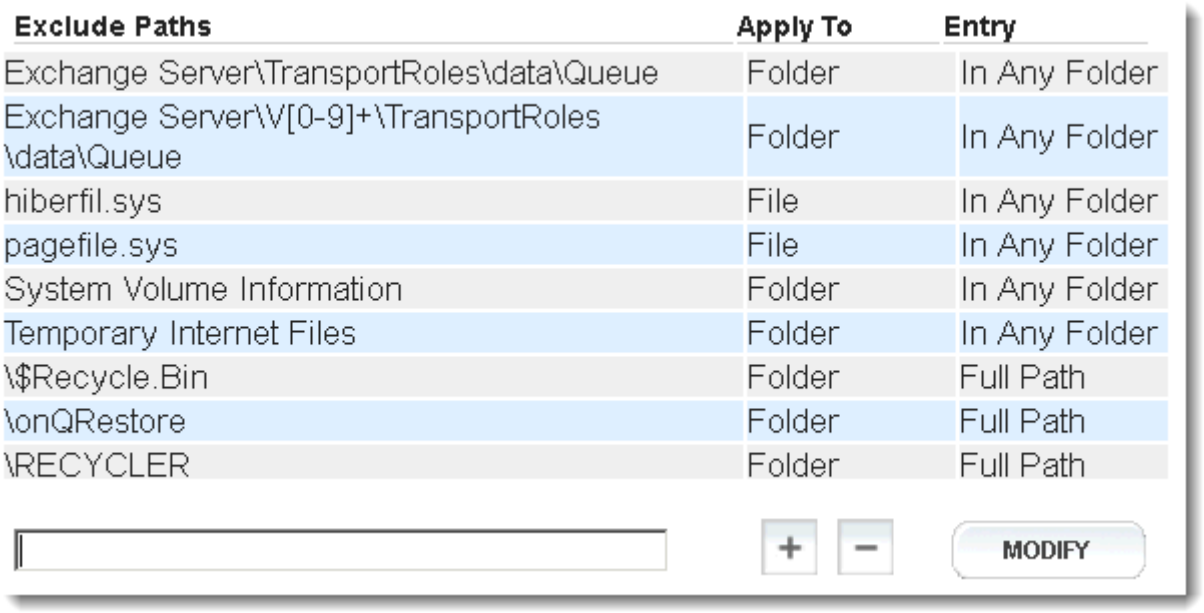
      & (ampersand)

   b. Select an **Apply To:** radio button that represents the path type. The path can be the folder or file or both.

Note: For a CIFS mount, you do not need to specify the full path of the mounted share volume. Simply type the drive letter and folder name: `P:\applications` where `applications` is the folder that you want to exclude.

c. Select an **Entry:** radio button to indicate that you'd like to exclude the file or folder wherever it is found or if you want to exclude only that specific full path.

In the following example, the contents listed will not be backed up. These exclude paths represent some of the paths that are listed in the *default* Backup Exclude List defined by Quorum.



Exclude Paths	Apply To	Entry
Exchange Server\TransportRoles\data\Queue	Folder	In Any Folder
Exchange Server\V[0-9]+\TransportRoles\data\Queue	Folder	In Any Folder
hiberfil.sys	File	In Any Folder
pagefile.sys	File	In Any Folder
System Volume Information	Folder	In Any Folder
Temporary Internet Files	Folder	In Any Folder
\$Recycle.Bin	Folder	Full Path
\onQRestore	Folder	Full Path
\RECYCLER	Folder	Full Path

6. Click the plus (**+**) button.
7. Do one of the following:
   - Click **Save** - onQ Portal saves your exclude list as shown in the list, then returns you to the previous page. The **Backup Exclude List:** button is labeled **Custom**. You cannot save the list if the total length of all paths exceeds 2,000 characters. Also, each line cannot exceed 512 characters.
   - Click **Default** - onQ Portal replaces the entire list with the default list (also known as *factory default*), then returns you to the previous page. The **Backup Exclude List:** button is labeled **Default.**
   - Click **Revert** - onQ Portal discards your changes, then returns you to the previous page. The **Backup Exclude List:** button does not change.

9.4 Edit backup include list

This feature (via the **File Backup Only** parameter) feature can only be enabled by Quorum Support and is typically used only for file servers.

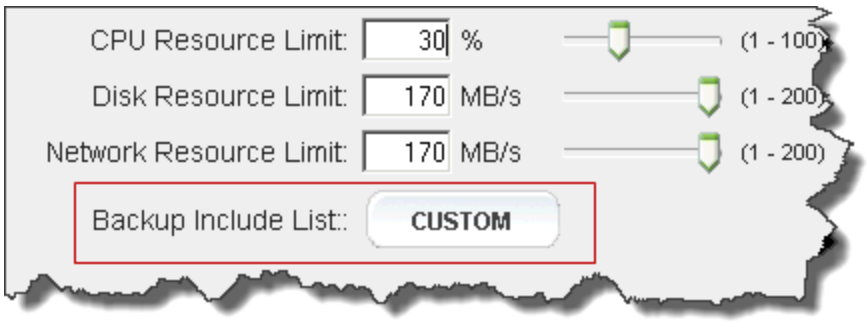
A backup include list is a user-defined set of folders/files that need to be backed up. The purpose of the list is to improve the efficiency of the backup process by being able to specify only the necessary files without having to create a large exclude list. For each protected node one can define a list of files or folders.

For an include list, Quorum recognizes file- and folder-names expressed as IEEE POSIX Basic Regular Expressions (BRE). You will find a synopsis of BRE syntax at http://en.wikipedia.org/wiki/Regular_expression.

To add a file or folder to the include list:

This procedure assumes that you're modifying the include list for an existing protected node. However, you can perform this procedure from either the **Add Protected Node** page or the **Modify Protected Node** page. Include lists are specific to the protected node's operating system.

1. Log on to the HA Appliance's onQ Portal.
2. Click the **PROTECTION CONFIG** tab.
3. Select the node, then **Modify**.
4. Select **File Only Backup** > **Yes** radio button > **ADVANCED** button.
5. In the **Backup List** field, click **Custom.**



6. Enter a path/folder/file and click plus (+) to add the entry to the include list.

7. Do one of the following:
   - Click **Save** - onQ Portal saves your exclude list as shown in the list, then returns you to the previous page. The **Backup Exclude List:** button is labeled **Custom**.
   - Click **Revert** - onQ Portal discards your changes, then returns you to the previous page. The **Backup Exclude List:** button does not change.

9.5 Initiate immediate backups

You can trigger an immediate backup of protected node, regardless of when the next backup is scheduled. You can restore from a backup at any time.

To initiate a backup:

1. Log on to the HA Appliance's onQ Portal.
2. Go to **DASHBOARD** tab > **PROTECTED NODES** page.
3. Unlock the page.
4. In the **Backup Status** column, click the carat icon for the protected node that you want to back up.
5. When prompted, click **Start**.
   - The carat changes to green and moves left to right until the backup completes.



   - After completion, the carat moves to its upright position and the new backup timestamp posts. In addition, the snapshot is saved to the repository, and you can retrieve it at any time.

Related Topics

Initial Backup
Schedule backups
Initialize the DR Repository
Resize protected node's vdisk

9.6 Export or import Repository

An export is most commonly used on the HA Appliance in conjunction with a USB device to:

- Initialize (“seed”) the DR Appliance, whether on-premise or Hybrid Cloud. To do this, simply use the export function on the HA Appliance and import function on the DR Appliance.
- Archive a complete repository.
- Move the repository from one onQ Appliance to another (a replacement, for example) when the data is too large for a LAN/WAN synchronization. See (Workflow) Fail back DR to HA Appliance.

In order to first synchronize the DR Appliance to an HA Appliance, a large volume of data must be transferred.

If the volume of data is high relative to the network bandwidth available between the onQ Appliances, it might be more efficient to export onQ’s deduplicated snapshot storage (aka “repository”) to a removable drive on the HA Appliance and then import it to the DR Appliance.

To export a repository:

1. Verify that the USB device has enough space to accommodate a copy of the repository.
2. Enable the export/import feature. Simply insert the device on the HA Appliance. onQ automatically enables the export/import function.
   - The **Export** button appears when an EXT4 formatted removable storage device is detected with adequate space for the local repository.
   - The **Import** button appears when a removable storage device is detected, and the device has a valid local repository.
3. Log on to the HA Appliance’s onQ Portal.
4. Stop protection on the HA Appliance.

Warning: Do not start protection until after you import the HA Appliance’s repository into to the DR Appliance. The procedure below walks you through this process.

5. (Important!) Allow time for any in-process backups/transfers to complete. Mouse-over the RN status for current state information.

   Alternatively, you can stop the backups as outlined in "Stop in-progress backups" on page 220.

6. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **REPOSITORY** page.

   There might be a short delay as onQ is calculating the size of the local repository and analyzing any removable storage devices.

7. Click **Export**.

   A progress bar appears, indicating that the export is in progress.

8. After the export completes, follow the procedure below to import the repository.

To import a repository:

An import permanently deletes the local repository and cannot be undone.

1. Enable the export/import feature. Simply insert the device on the DR Appliance. onQ automatically enables the export/import function.
2. Log on to the DR Appliance's onQ Portal.
3. Stop protection on the DR Appliance.
4. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **REPOSITORY** page.
5. Click **Import**.
6. Wait for import to complete.
7. Start protection on HA Appliance.
8. Start protection on DR Appliance.

Related Topics

Initial Backup
Initialize the DR Repository

9.7 Schedule backups

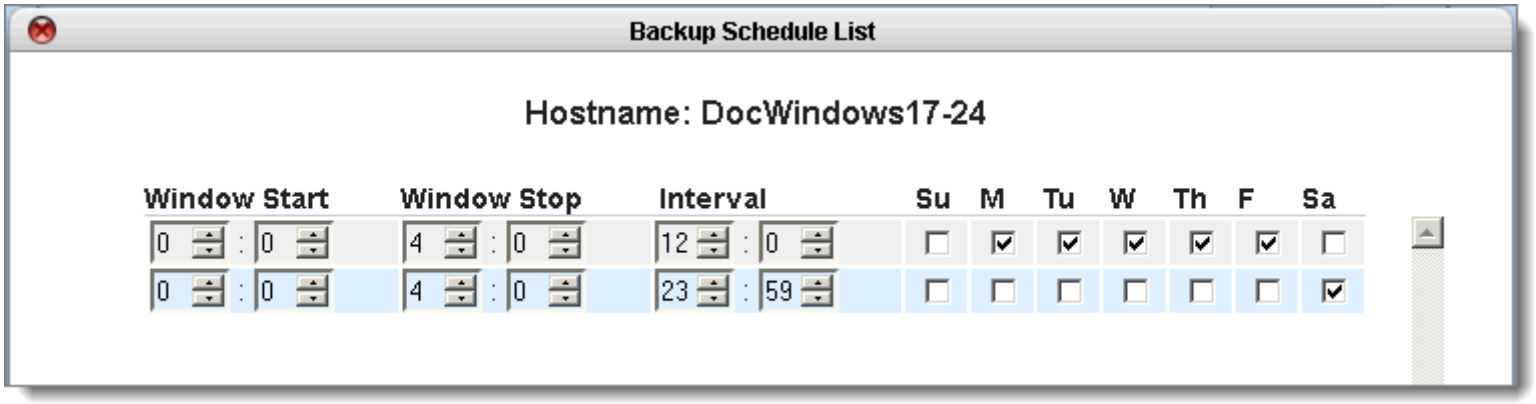
onQ performs an initial backup of your PN when you first start protection on your nodes. For all subsequent backups, you can instruct onQ when and how often to perform those backups. Each protected node has its own backup schedule with backups to occur at varying times, days of the week, and frequencies. Such granular scheduling is essential for those networks with low bandwidth.

To create or change the backup schedule:

This schedule is automatically transferred to the DR Appliance. In the event of a disaster and if the DR Appliance takes on the HA role, the same backup schedule continues.

1. Log on to the HA Appliance's onQ Portal.
2. Go to **PROTECTION CONFIG** tab, select the PN, then click the **MODIFY** button.
3. In the **Backup Schedule** field, click the **SCHEDULE** button.
4. Configure your schedule. You can create as many schedules as you want. The scheduler prevents you from creating schedule overlaps. If you want to return the default schedule, click **CLEAR**.

   For example, let's assume you have very little network bandwidth and peak hours of operation are M-F 5am to 11pm and business activity is light on Saturdays with minimal changes and no activity on Sundays. Simply set the schedule as follows:



Window Start: Based on a 24-hr clock, select the time of day to start the backup. Choose a time when the network is the least busy. The initial backup can take a considerable amount of time. Subsequent backups record only changes that have occurred since the preceding backup and so will be relatively brief.

Window Stop: Based on a 24-hr clock, select the time of day to stop the initiation of new backups. Any backups currently running at the time of the stop period will be allowed to complete. As outlined in Stop in-progress backups, you can stop these backups.

Interval. The interval (in minutes/hours) between the start of the data backup and the next data backup. Basically, the interval represents the frequency of backups. Valid values: hour, `0-23`; minute, `0-59`. As with any backup solution, this interval represents your potential data loss in the event of a failure (aka RPO).The shortest interval is 15 minutes, which is impractical for an onQ that's protecting a large number of PNs, possibly resulting in queued backups as outlined in "Stop in-progress backups" on page 220. If you are not sure what interval to specify, a 4-hr interval is common in enterprises with high network bandwidth and 24-hr interval in low network bandwidth.

Days: Select the check box for the day(s) of the week for which you'd like this policy to apply. If you don't have business activity (for example, Sunday) on a given day, clear that check box.

That's it!

5. (Optional) Monitor these backups. Go to "Monitor protected nodes" on page 358 and "Monitor backups" on page 369. If you receive any email alerts, follow the instructions in Backup Alerts.

9.8 Change Backup Retention Policy

You must tell onQ how long to retain backups for a given PN. Each PN has a unique backup retention period for both the HA Appliance and the DR Appliance.

In the event that onQ needs to free up disk space, it will expire backups that agree with your retention policy. In the event that onQ is low on disk space, you will receive one or more alerts outlined in onQ Disk Space Alerts. If onQ is unable to expire backups because doing so conflicts with your retention

policy, you might need to reduce the number of days that you specified in the node's retention policy.

To change the backup retention period:

1. Log on to the onQ Appliance's onQ Portal.
2. Click the **PROTECTION CONFIG** tab.

   The PROTECTION CONFIG tab on a DR Appliance is slightly different from its HA counterpart. The DR's tab shows the HA Appliance's backup retention time as well as its own backup retention time:



3. Select the node that you want to modify, then click **MODIFY**.

   The Modify a Protected Node page appears.
4. In the **Backup Retention Time** field, specify (in *days)* how long you want onQ to retain backups, then **SAVE**.

Related Topics

Delete snapshots
Schedule backups

9.9 Delete snapshots

You cannot delete snapshots for an existing PN. However, if you deleted a PN, because you either removed the PN from production or you moved the PN to another HA Appliance, you might want to delete all the orphan data associated with that PN. Such orphan data includes:

- PN snapshots
- RN
- RN snapshots
- Self test schedules

However, in the latter case, don't delete the orphan data without first confirming that backups for the PN are working on the new HA Appliance. Moreover, to keep the onQ Appliances in sync, Quorum recommends that you perform this procedure on the HA Appliance's DR Appliances and DR Mirrors.

To delete orphan data:

1. Log on to the HA Appliance's onQ Portal.
2. Go to **PROTECTION CONFIG** tab > **FIND ORPHANS** button.

   The onQ Portal searches the Repository for orphan data. If onQ Portal doesn't find any orphan data, the onQ Portal returns a **`No data found that does not belong to a configured PN`** message. Otherwise, the onQ Portal prompts you to purge the orphan data that the onQ Portal did find.
3. Select the check box for the PN whose orphan snapshots you want to delete, then click **YES**.

Warning: You cannot abort a purge after it begins.

   A progress dialog appears, highlighting the number of files waiting to be deleted. This process can take several minutes per snapshot.
4. After the purge completes, repeat Step 2 to verify that the onQ Portal removed all orphan snapshots.
5. (Recommended) Repeat this process on each DR Appliance and DR Mirror so that all onQ Appliances are in sync.

9.10 Stop in-progress backups

Whether you use agent-less or agent-based enrollment, onQ can perform

three PN backups simultaneously. However, with agent-less enrollment, your agent-less PNs must reside on different ESX/ESXi hosts.

Even when you schedule backups during non-working hours, backups can occur during unscheduled times (working hours) due to queued snapshots. Queued snapshots can occur if onQ is overly busy backing up other PNs or if the PNs have short backup intervals.

If an unwanted backup is in progress during unscheduled times, you might want to stop that backup from continuing if it is affecting performance.

To stop a backup:

1. Log on to the HA Appliance's onQ Portal.
2. Go to **DASHBOARD** tab > **PROTECTED NODES** page.
3. Unlock the page.
4. Locate the PN with a backup in progress as indicated by the **Backup Status** column (see Monitor backups).
5. Click on the Backup Status button. The **Backup Status in progress** dialog appears.
6. Click **STOP BACKUP**.



The backup will automatically run at the next scheduled interval as displayed in the **Next Scheduled Backup** column.

7. (Optional) Consider increasing your backup interval to prevent queuing of backups.

9.11 Retrieve objects from Repository

To browse or retrieve files/folders in the Repository, you have two choices.

- Indirectly, using FLR. Copy the files to another Windows server, then browse.
- Directly, using WSR. Simply browse *any object* in the Repository directly or map to a shared location on a Windows server or client.

Note: With WSR, you can use a 3rd party exchange recovery tool and open the exchange database directly from the Repository. In doing so you can retrieve any object, even a single email message within couple minutes and from any prior snapshot.

9.11.1 Perform WSR Restore

A Windows Share Restore (WSR) restore involves mapping from a drive on a Windows server/client to a shared location generated by the onQ. This shared location will contain all the snapshots for a given PN.

For security purposes, you can only mount the share from one location; besides, one share is recommended because browsing the repository can be resource intensive.

After you create the share, you can browse any of the snapshots (backups) listed for viewing, reading, and copying any folders/files as needed. However, a few things to note:

- You cannot modify or delete folders/files within a snapshot.
- You cannot delete folders/files from the repository.
- You cannot copy new folders/files into the repository.

To restore Exchange Server data and mailboxes, consider using Kroll's

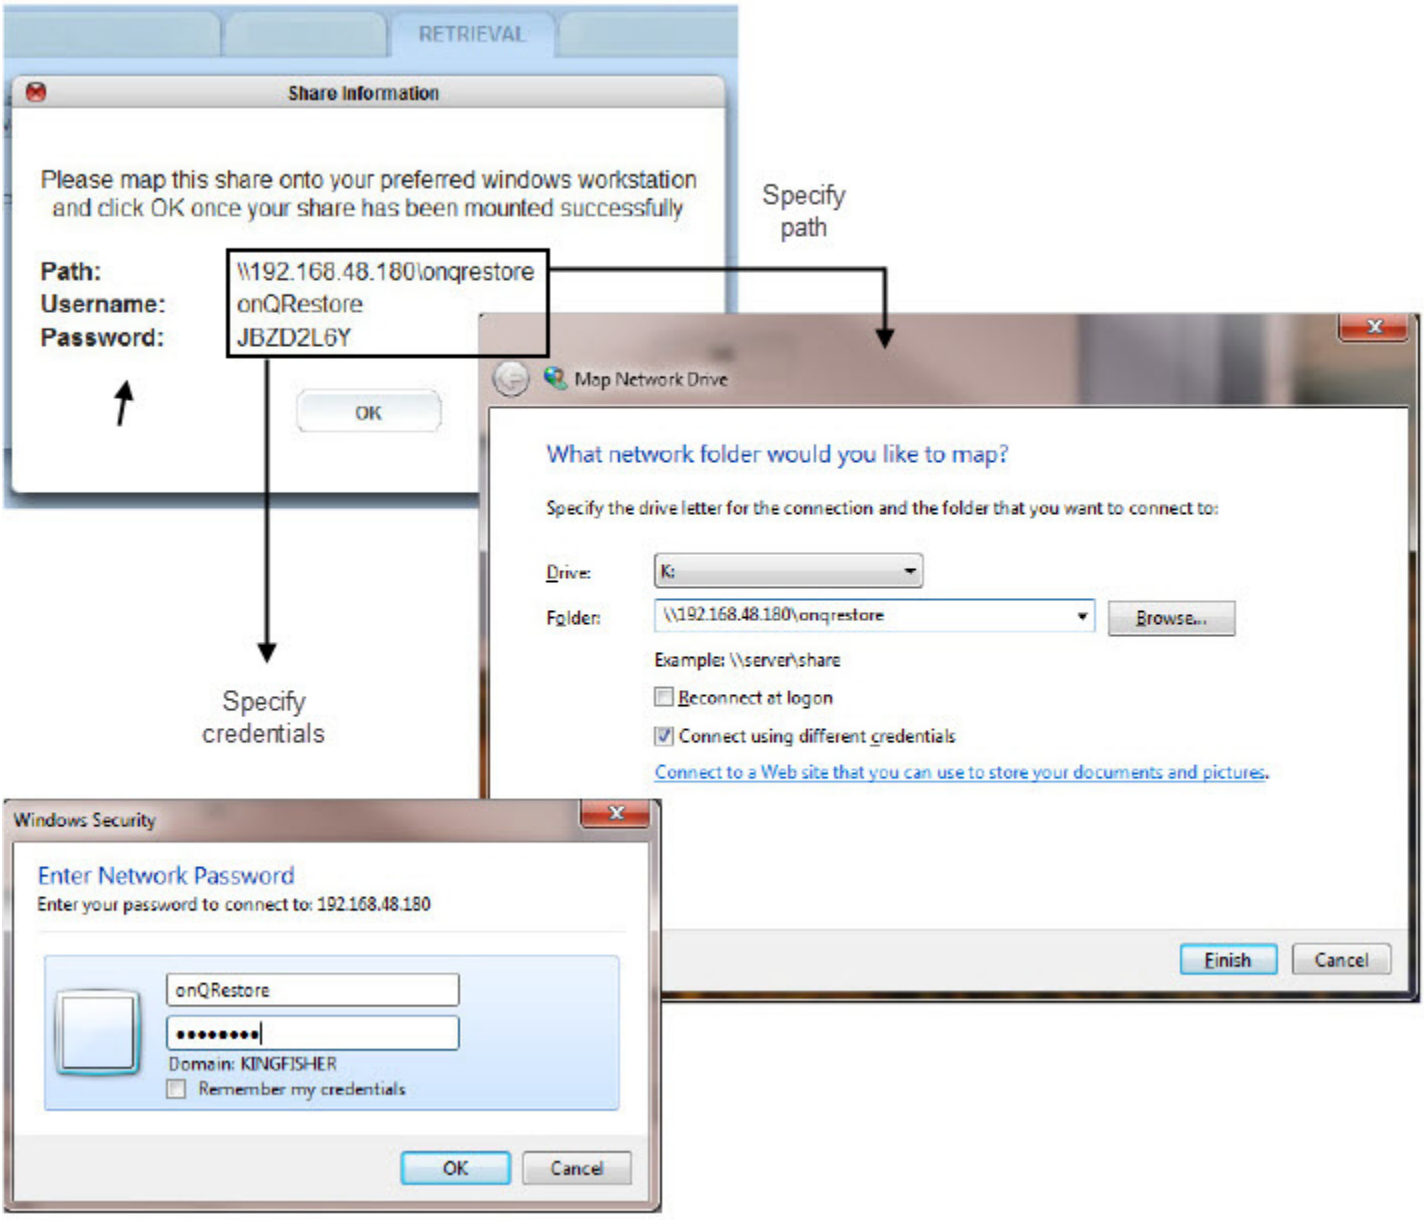
Ontrack PowerControls.

To retrieve a snapshot, using WSR:

1. Log on to the Appliance's onQ Portal.
2. Click on the **RESTORE** tab > **WINDOWS SHARE** page. onQ displays an empty page if no share was activated.
3. Create a share, then map to this share:
   a. Click on the **ACTIVATE SHARE** button.
   b. From the Active Windows Share Session popup and in the drop-down box, select the desired protected node from which you want to restore, then **START**.

      When a protected node appears in the list, that PN has data in the Repository, including any unprotected nodes that were previously protected.
   c. Using the share path, username, and password that appears in the Share Information popup, map your Windows server/client to this share, then **OK**.

      The share expires after 30 minutes of no activity. Also, if the share is not mapped on to a Windows server/client before you exit the popup, both `Client` and `Connection Start Time` values are blank.



4. Launch Windows Explorer to display the snapshots available for the PN.
5. From here, copy the data to any restore location.

To delete a share:

You might want to delete share if you:

- You completed your restore or no longer need to retrieve objects from a snapshot.
- You forget to record the credentials for the share location. For security purposes, you cannot retrieve this information, and so you must delete and recreate the share.
  1. Log on to the Appliance's onQ Portal.
  2. Click on the **RESTORE** tab > **WINDOWS SHARE** page.
  3. Select the share from the list, then **STOP SHARE**.
  4. In the Stop Windows Share Session popup, click **YES**.

9.11.2 Perform File-level Restore

onQ provides a way for you to retrieve any combination of files and folders or entire drives from existing snapshots of any protected node. This process is known as a file-level restore (FLR).

If, however, your protected node's server failed, you need to perform a bare metal restore (BMR) as outlined in "Restore PNs Using QUARK" on page 315. You can choose to *replace* the current file, folder, or drive, or to *store* it on any protected node or standalone server so long as it is reachable from your network.

Note: FLR is supported for both Windows and Linux-based protected nodes, though the example below demonstrates a recovery using a Windows-based protected node.

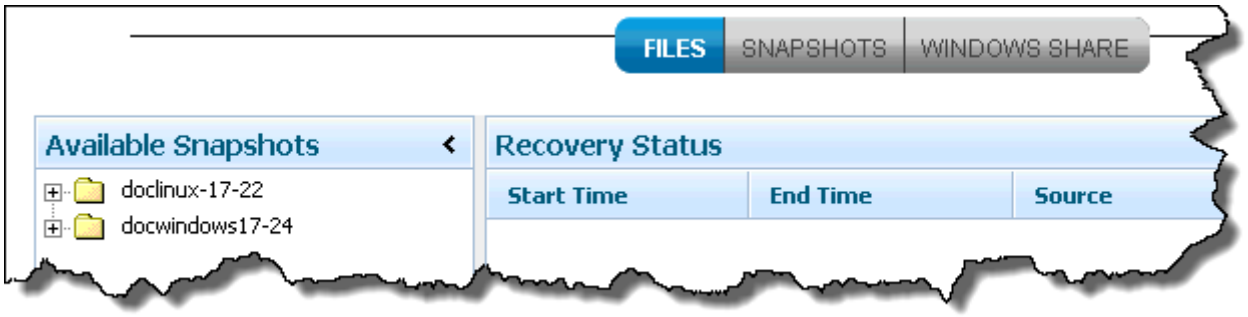
To retrieve a snapshot, using FLR:

If you have a 6.x RHEL PN, ensure that you have the `nc` utility installed as outlined in Step 6 of (Agent-based Linux PNs) Enroll protected nodes.

Before you begin, review the limitations outlined in "File-level Restore Limitations" in onQ Release Notes.

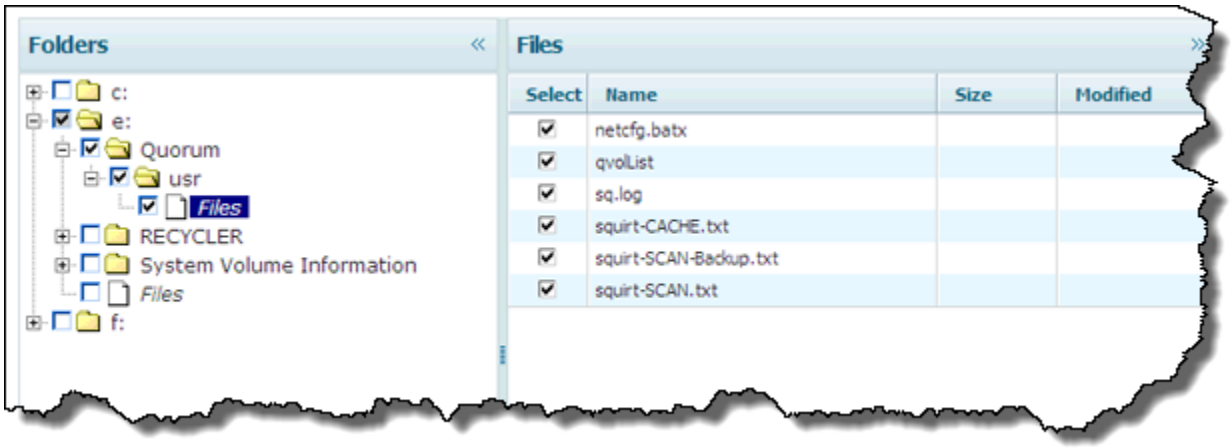
1. Log on to the Appliance's onQ Portal.
2. Click on the **RESTORE** tab > **FILES** page.

A page similar to the following appears:



3. In the **Available Snapshots** pane, expand the folder for the protected node.
4. Double-click on the snapshot that contains the data you want to restore.

   Another page, labeled **Step 1**, appears:



5. Specify the files you want to retrieve. Expand the hierarchy on the left, select the desired folder or files, then click **NEXT**:
   - For each folder, select the check box to select all of its contents.
   - Click ***Files*** to display the individual files in the content pane to the right. Then, in the content pane, you can select the individual files that you want.

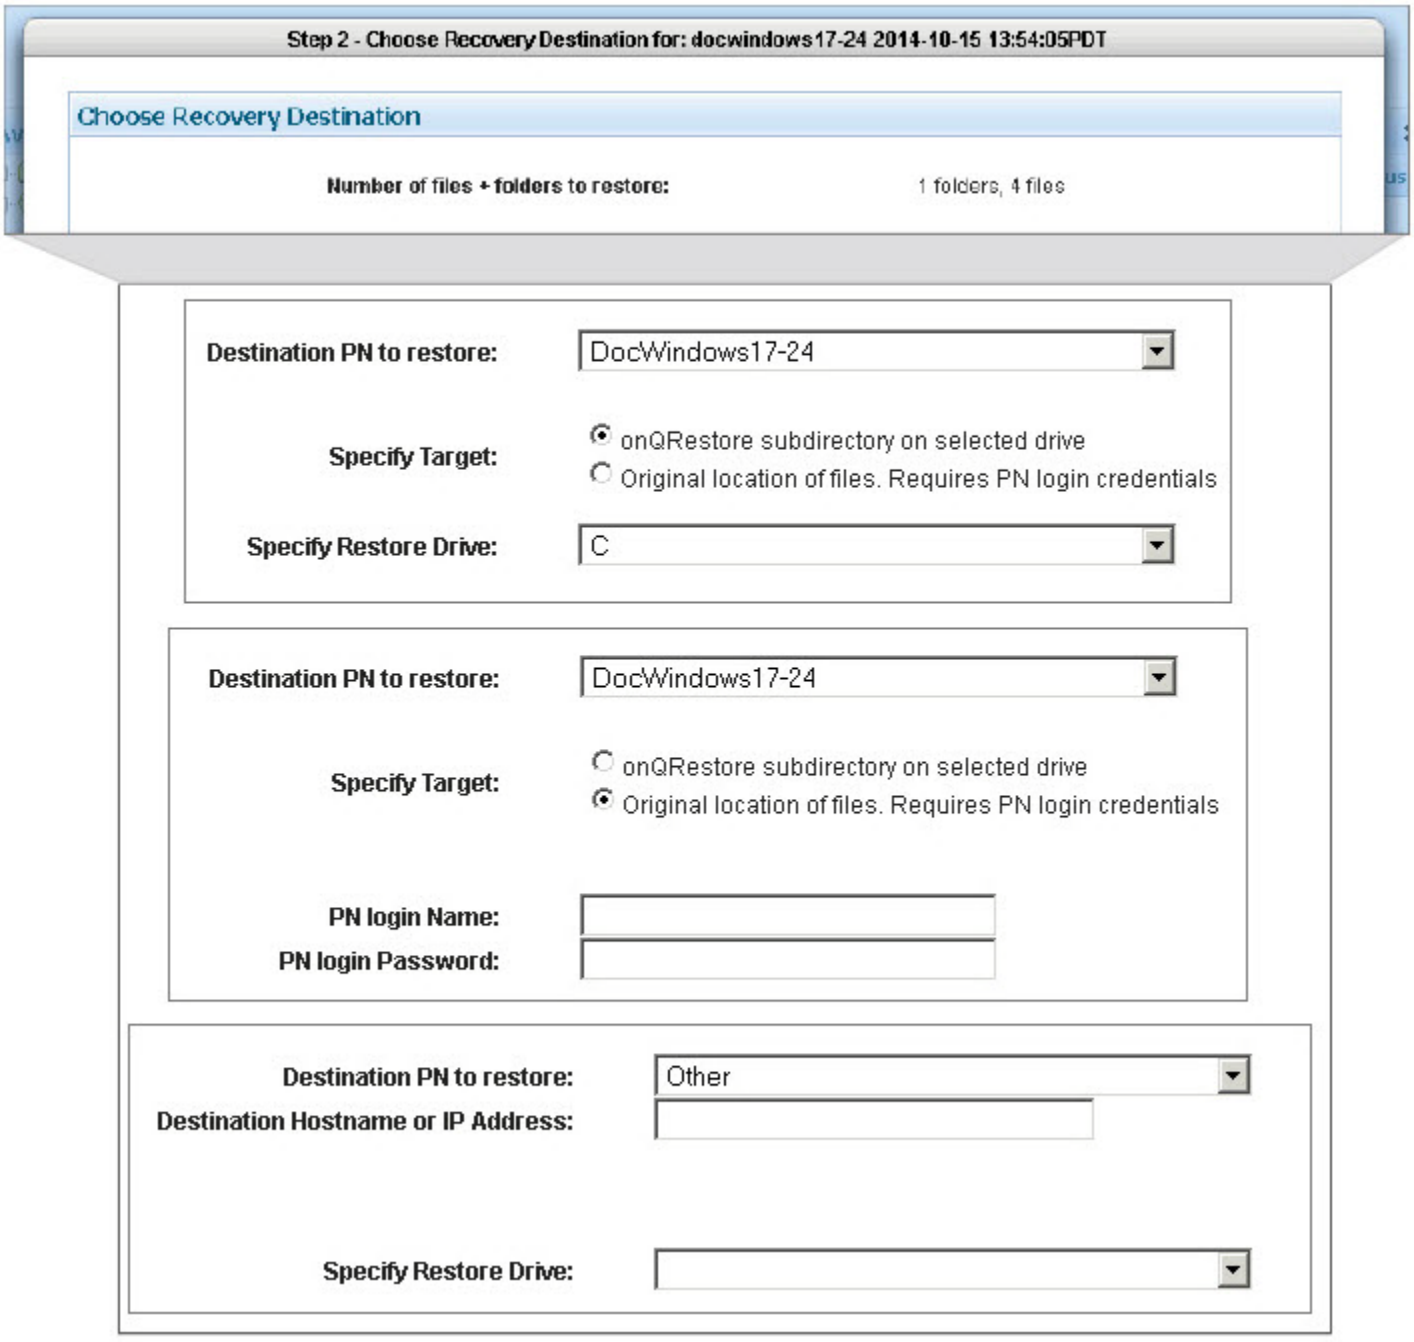
   Another page, labeled **Step 2**, appears:

6. Before you proceed with the restore, decide on the host to which you want to restore the files.
   - **Do you want to restore to a *Protected Node***? You can restore to the same or different Protected Node.
   - **Do you want to restore to a *Standalone Server***? Do so if you'd like to inspect the files before you restore a Protected Node with them.
7. Do the following:
   a. In the **Destination PN to restore** drop-down, specify the recovery destination—the host where you want to store the retrieved files. The destination can be one of the following:

      <***PNHostname***>. This destination is a PN, either the same PN or a different PN. Choose this option if you chose "Protected Node" in Step 6. To ensure that you restore your files correctly, only PNs that share the same operating system as the host that owns the snapshot that you want to restore appear in the list. If applicable, in the **Specify Restore Drive** drop-down, choose a drive to which you want to save the files

      Other. This destination is a host that is not a PN. Choose this option if you chose "Standalone Server" in Step 6. If you have an agent-less PN, this is the only supported option for FLR. When you select **Other**, a text box appears labeled **Destination hostname or IP address**. Type the address of your desired destination

server. In the **Specify Restore Drive** drop-down, choose a drive to which you want to save the files. Skip to Step c.



b. If applicable, in the **Specify Target** field, choose a target. A target can be one of the following: (1)

- **onQ Restore subdirectory on selected drive**. The files will

automatically be saved to a default restore location on the drive that you select and that distinguishes the FLR from multiple different snapshots and maintains the relation of the files to a given snapshot:

(Windows) – *driveLetter*`:\onQRestore\`*pnName AND snapShot timeStamp*.

(Linux) – *mountPoint*`/onQRestore/`*pnName/snapShot timeStamp*. `/onQRestore`... if you chose `/`or `/boot/onQRestore`... if you chose `/boot`.

- **Original location of files**. Choose this option if you want to overwrite the existing files. There are limitations to this option (see "File-level Restore Limitations" in onQ Release Notes). Skip to Step 8.

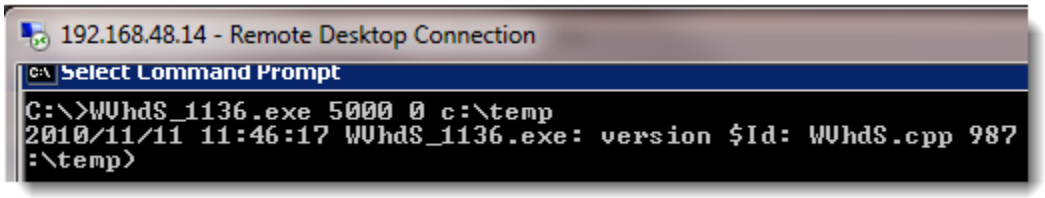
c. If you want to restore to a Standalone Server, which, unlike a restore to a PN, does not already have the onQ Service installed, do one of the following to open up a communication "channel" between the onQ Appliance and the Standalone Server:

- **(Windows/Automatic Process)** Type the PN's credentials, then wait for `WVhds_<build>.exe` to automatically validate those credentials. These credentials are not stored on the onQ in any database or log file.
- **(Windows/Manual Process)** Download and execute `WVhdS_<build>.exe` on the remote host as the example below shows. Retrieve this file from **APPLIANCE CONFIG** tab > **ADVANCED** button > **DOWNLOADS** subtab. From a DOS prompt, run the following command:

```
# WVhdS_version.exe 5000 0 restore-location
```

For example:

(where `c:\temp` in the following example is the restore location)



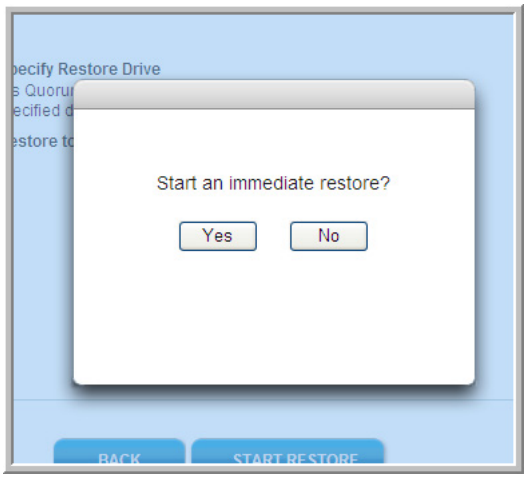
- **(Linux/Manual Process)** Download and execute `onQFLRSelf.sh`, which automatically untars a file to the recipient that you define, on the remote host. Retrieve this file from

APPLIANCE CONFIG tab > **ADVANCED** button > **DOWNLOADS** subtab.

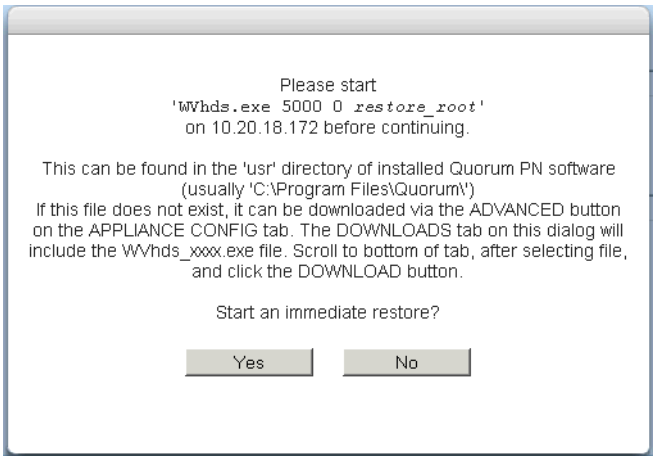
```
# cd <restore-location>
# wget http://<OnQ-IP-address>/onQFLRSelf.sh
# sh onQFLRSelf.sh
```

8. Begin the restoration. Click **START RESTORE**, then click **Yes** to confirm:

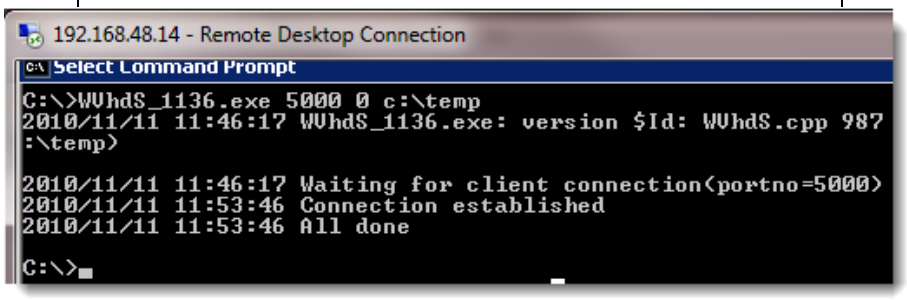
Note: Protection must be ON in order for the **START RESTORE** button to appear.

If you chose to restore to a PN:	

If you chose to restore to a Stand-alone Server:



If restore completes, output indicates `All done`:



Afterward, the onQ Portal returns you to the main **RESTORE** page.

9. Verify that your restore completed successfully:

   a. Observe Snapshot status (**RESTORE** tab > **FILES** page > **Recovery Status** pane > **Status** column):
      - The left pane (**Available Snapshots**) - shows available snapshots, as folders, listed in a hierarchical structure by node, date, and time.
      - The right pane (**Recovery Status**) - shows your previous efforts at retrieving files from snapshots. The carat to the right

symbolizes the result of each restore attempt.

Status	**Completed OK**  carat is green and points upward	Restore finished successfully. If you roll your cursor over the icon, you see the message.
	Incomplete(NetworkErr)  carat is red and points downward	Restore failed. If you roll your cursor over the icon, you see a reason for the failure. This error can occur if the Netcat (`nc`) utility is not installed on the target Linux node as outlined in Step 6 of (Agent-based Linux PNs) Enroll protected nodes. Also, if this utility is missing, your var/log/messages indicate: `/opt/quorum/bin/./FLR.sh : line 44: /usr/bin/nc: No such file or directory`.
	Incomplete Client: Unable to connect to FLR server  carat is red and points downward	Restore failed. If you roll your cursor over the icon, you see a reason for the failure.
	Incomplete(ReplyError)  carat is red and points downward	Restore failed. If you roll your cursor over the icon, you see a reason for the failure.
	Completed with errors  carat is yellow and points upward	Typically this error occurs when your target option is `Original location of files`. To view the errors, open the `quorum-wvhds.log` file on target.

Incomplete: Unexpected Disconnect	Restore was interrupted, probably due to a timeout.

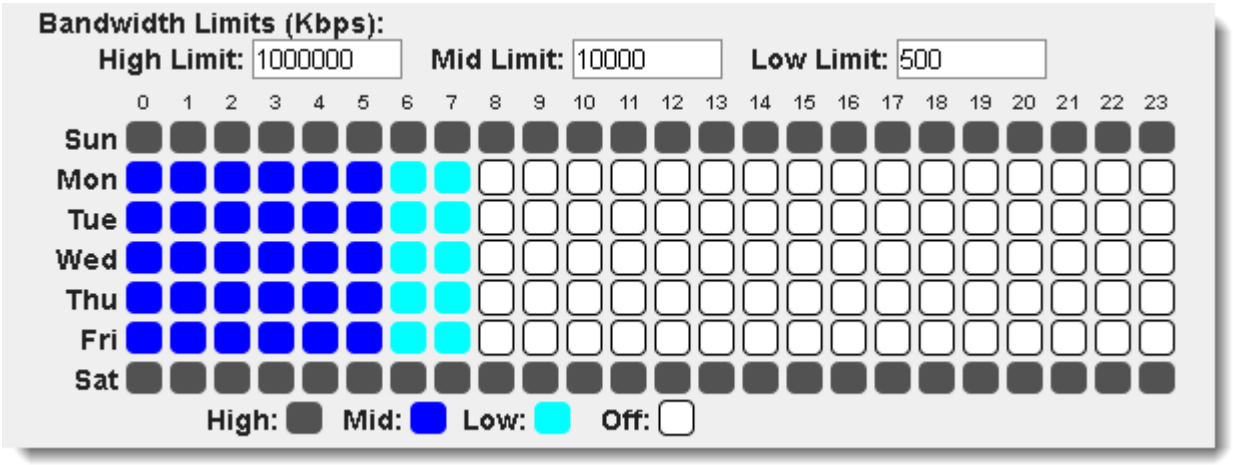


- (Standalone Server) Verify that the restored files appear in the expected destination.

9.12 Adjust bandwidth throttling

Bandwidth throttling on the WAN connection between the HA Appliance and the DR Appliance is set by default to the maximum bandwidth available. If you have a DR Mirror, you can also adjust bandwidth throttling between the DR Appliance and the DR Mirror.

Bandwidth throttling can prevent a single transfer operation from swamping the network, node crashes, or other processing bottlenecks. You might need to adjust bandwidth throttling downward if the replication process places too much of a burden on your network.



To adjust bandwidth throttling between HA Appliance and DR Appliance:

1. Log on to the HA Appliance's onQ Portal.
2. Go to **APPLIANCE CONFIG** tab > **onQ (REMOTE)** page, then **MODIFY**.
3. Change the value (Kbps) in the **High Limit**, **Mid Limit**, and **Low Limit** fields.
4. Use the grid by clicking on the boxes to cycle through the bandwidth settings to select the limit for each day and hour, then **SAVE**.

To adjust bandwidth throttling between DR Appliance and DR Mirror:

1. Log on to the DR Appliance's onQ Portal.

2. Go to **APPLIANCE CONFIG** tab > **onQ (REMOTE)** page, then **MODIFY**.
3. Change the value (Kbps) in the **High Limit**, **Mid Limit**, and **Low Limit** fields.
4. Use the grid by clicking on the boxes to cycle through the bandwidth settings to select the limit for each day and hour, then **SAVE**.

Related Topics

Limit resource utilization on PNs

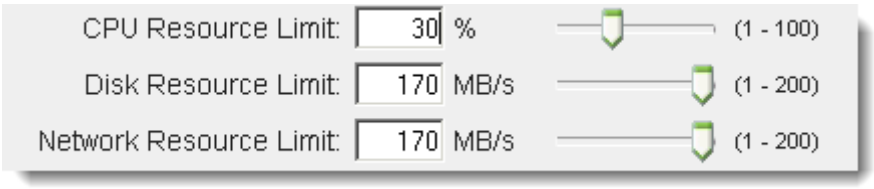
9.13 Limit resource utilization on PNs

When onQ Manager performs certain tasks, such as backing up protected nodes, it uses resources on those hosts. onQ Manager is always aware of the resource utilization on your hosts.

Your HA Appliance is pre-configured to use normal resource utilization by default, but you can further limit the CPU, disk, and network resources that onQ Manager uses. These controls can provide you assurance that your end users receive the highest priority during peak application usage.

To set resource limitations:

1. Log on to the HA Appliance's onQ Portal.
2. Stop protection.
3. From the **PROTECTION CONFIG** tab, select the protected node from the table, then **MODIFY** > **ADVANCED**.
4. Specify the resource limits using either the slider or the text boxes provided. Go more information on these limits, go to Protected Node Parameters.



Related Topics

Adjust bandwidth throttling

9.14 Run custom backup scripts

onQ supports scripts to enable you to quiesce an application or database for backup without manual intervention.

(Agent-based PNs) To enable custom backup scripts:

To quiesce an application or database for backup without manual intervention, simply use `pre-snap` and `post-snap` scripts:

- For Windows PNs, you can define batch file scripts that run automatically before and after the HA Appliance creates a VSS snapshot.
- For Linux PNs, you can define shell scripts that run automatically before and after the HA Appliance creates a backup. If, for example, your pre-snap script shuts down your database so as to back up the database, the database remains shutdown until the post-snap script executes the startup process.

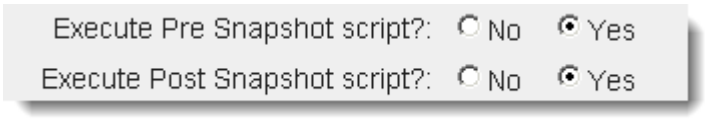
You can do so on a per-protected-node basis by saving your scripts to the PN, then registering them with the HA Appliance.

If you want to run these same scripts on all protected nodes, you must repeat the procedure below for each protected node.

1. Create two scripts with the following names, specifying the commands that you want executed.
   - **`pre_snap.bat`** and **`post_snap.bat`**. For Windows, click here to download `.bat` placeholder files to quiesce a DB2 database. In this example, each of the batch files in turn executes `.clp` scripts. These `.clp` scripts quiesce the DB2 database before the snapshot starts and again after the snapshot completes. After you create/install these scripts, they are not overwritten during a PN upgrade.
   - **`pre_snap.sh`** and **`post_snap.sh`**: For Linux, these placeholder files were saved to the destination you chose during the installation. The default location is `/opt/quorum/bin`. These scripts are not overwritten during a PN upgrade.

2. Log on to the PN and save these scripts to:
   - (Windows) `\Program Files\Quorum\usr` folder.
   - (Linux) the destination you chose during the installation. Usually, `/opt/quorum/bin/` is the preferred directory.
3. Log on to the HA Appliance's onQ Portal.
4. Stop protection.
5. From the **PROTECTION CONFIG** tab, select the protected node from the table, then **MODIFY** > **ADVANCED**.
6. Enable the pre and post snapshot scripts on that node:
   - Select the **Execute Pre Snapshot script?: Yes** radio button.
   - Select the **Execute Post Snapshot script?: Yes** radio button.

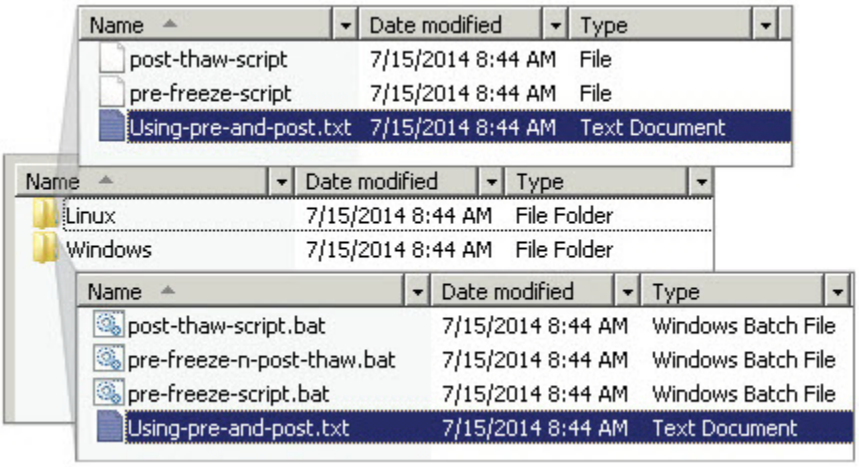
   The pre-snapshot script is `pre_snap`, and the post snapshot script is `post_snap`.



(Agent-less PNs) To enable custom backup scripts:

To quiesce an application or database for backup without manual intervention, simply use `pre-freeze` and `post-thaw` scripts:

1. Download and unzip the example scripts. There are two folders, one for Windows and another for Linux. Each folder contains sample scripts.



2. Create `pre-freeze-script.bat` and `post-thaw-script.bat` files, specify the commands that you want executed.
3. Do one of the following, depending on the PN's operating system:
   - (Windows) Save the `pre-freeze-script.bat` and the `post-thaw-script.bat` files under one of the following locations:
     - `C:\program files\vmware\vmwaretools\backupScript.d\` folder. When using this path, both the pre and the post scripts can be combined as in the `pre-freeze-n-post-thaw.bat` example script.
     - `C:\WINDOWS` folder on the Protected Node.
   - (Linux):
     - Save the `pre-freeze-script.bat` and the `post-thaw-script.bat` files under `/usr/sbin` of the Protected Node.
     - Make scripts executable (**chmod +x** <*script*>).
4. Initiate a backup of PN from onQ Portal.

Related Topics

Back up and restore Oracle 11g database on Linux

Back up and restore Oracle 10g+ database on Windows

9.15 (Agent-based PNs) Methods of Performing Incremental Backups

After onQ performs an initial snapshot of your protected nodes, onQ performs incremental updates from that point forward. onQ can employ one of two scan solutions to achieve this goal: Filter Driver or Non-Filter. The onQ Filter Driver is supported for agent-based PN enrollment only.

During enrollment, the Protect Me wizard installs the onQ Filter Driver, but it is disabled by default. If you want to choose this solution you must enable it.

onQ Filter Driver enables faster and more efficient incremental updates. The onQ Filter Driver is no different than the filter drivers that Windows backup products use. The major difference between the onQ Filter Driver and Non-Filter is that the onQ Filter Driver installs as a service on the protected node, using some system resources.

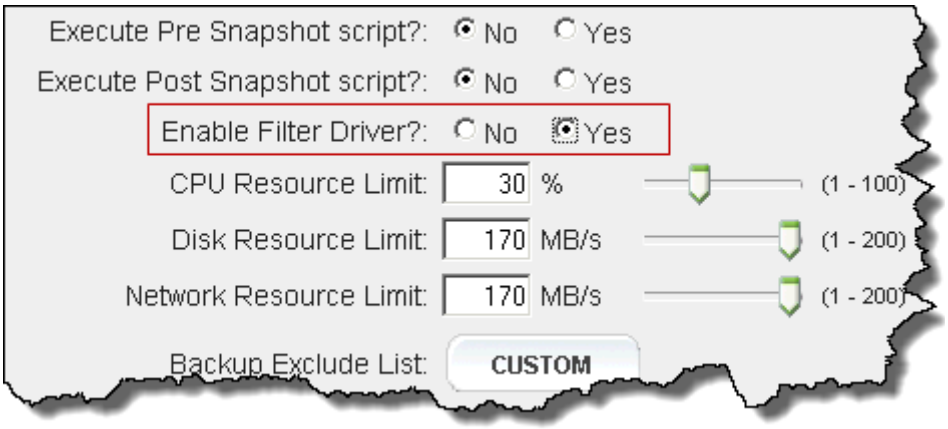
Though you don't need to know the "ins and outs" of each solution, you do need to understand when each performs optimally, so that you can choose

the right solution for your organization:

Circumstance	Recommendation	
	onQ Filter Driver (Optional)	Non-Filter (Default)
PN stores a large number (millions) of files	✓	
PN stores large files that change frequently	✓	
PN files don't change frequently		✓
PN stores a small number of files		✓
PN has resource constraints (memory, CPU, disk space)		✓

To enable the onQ Filter Driver:

1. Log on to the HA Appliance's onQ Portal.
2. Stop protection.
3. Go to **PROTECTION CONFIG** tab, select the protected node, then click **MODIFY** > **ADVANCED**.
4. Select the **Enable Filter Driver? Yes** radio button, then **SAVE**.



Related Topics

Configure automatic testing of RNs

9.16 Perform full rescan on PN

In most cases, you will not need to perform a full rescan of the node that you are protecting because the onQ handles all the changes based on the backup schedule that you configured. However, you might want to perform a rescan if:

- you are not comfortable with the delta
- a lot of changes occurred and you don't want to wait for the next scheduled backup

To initiate full rescan:

1. Log on to the HA Appliance's onQ Portal.
2. Go to **DASHBOARD** tab > **PROTECTED NODES** page.
3. Unlock the page.
4. In the Protected Node column, click on the node name button for the protected node that you want to rescan.
5. When asked if you want to **Force a Full Rescan on next backup?**, click **CONFIRM**.

9.17 Enable onQ to back up shared volumes

Aside from specifying the shared volume (mount point) at that time that you enroll the PN, in order for onQ to back up shared volumes, these shared drives need to include the required privileges. This requirement does not apply to fixed drives.

To set these privileges:

1. If your PN is a Win2k8 system, apply the hotfix 270891 as described in Microsoft's KB article 973278 to be able to back up NAS volumes:
   - 64-bit Win2K8
   - 32-bit Win2K8
2. Configure the onQ Service service to run as one of the following:
   - the same administrator that is currently logged on to the PN
   - Local administrator
   - Domain administrator
3. Ensure that the CIFS-enabled shares are mapped to the protected nodes with either Local administrator or a Domain administrator.
4. Specify the shared volume (mount point) at that time that you enroll the PN.

10

Security and Communications

- Test HA-to-DR Link
- Set up trust relationships
- Create Custom Networks for RNs
- Assign RNs to Networks
- Modify TCP Ports for HA-DR Communications
- Synchronize system time
- Secure connection to PNs

10.1 Test HA-to-DR Link

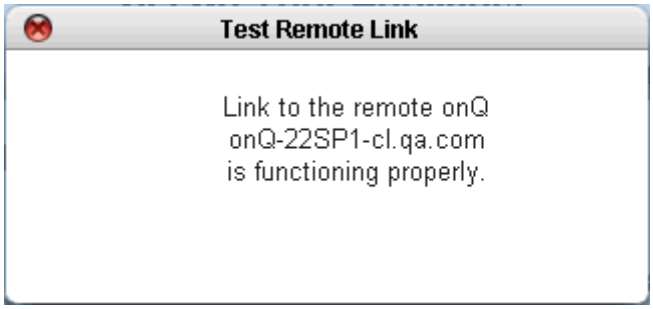
Use this procedure to:

- test the trust relationship between the Appliances.
- test inter-Appliance communications on port 81. This test does not test remote ssh communication.

To test your HA-to-DR link:

1. Go to **APPLIANCE CONFIG** tab > **onQ (REMOTE)** page.
2. Click **Test Remote Link**.

   If successful, you see a page similar to this:



Related Topics

Monitor DR Appliance
Network and Firewall Requirements

10.2 Set up trust relationships

Before your Appliances will work in production, you need to set up the trust relationship between the HA Appliance and the DR Appliance. Moreover, if you have a DR Mirror, you must set up the trust relationship between the DR Appliance and the DR Mirror.

Also, you need to re-establish this trust relationship between the HA Appliance and DR Appliance if you change their roles, which is required in the event of a disaster (see (Workflow) Fail over HA Appliance to DR Appliance).

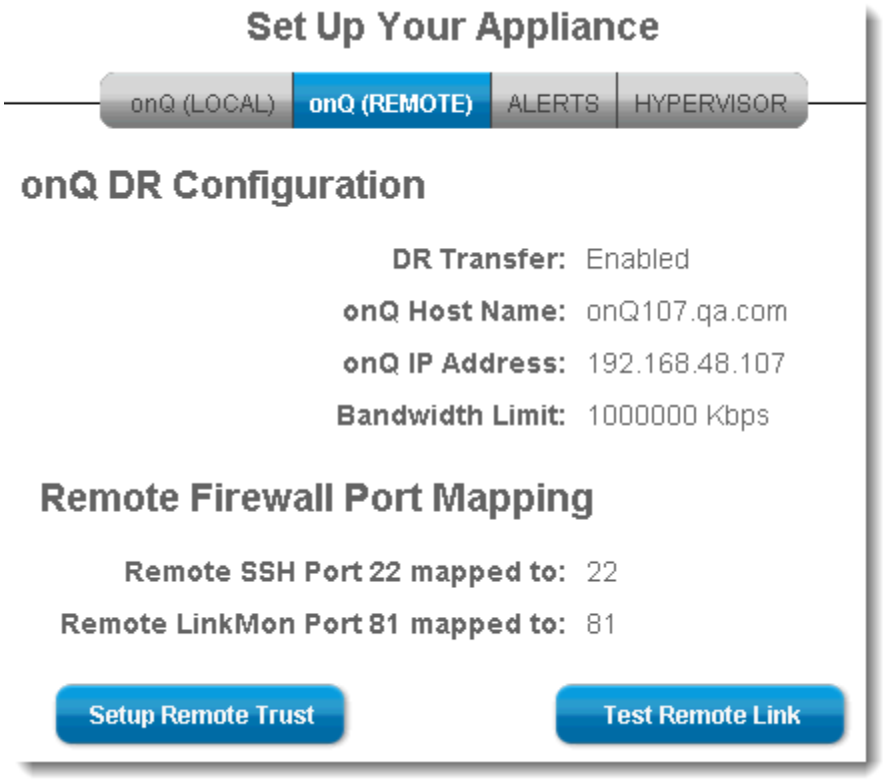
Therefore, just in case you have a disaster, it's best practice to set up the trust relationship in advance so that it's immediately available when you need it. As such, consider performing the following procedure on both the HA Appliance and, if you have a DR Mirror, the DR Appliance.

To set up the trust relationship between HA Appliance and DR Appliance:

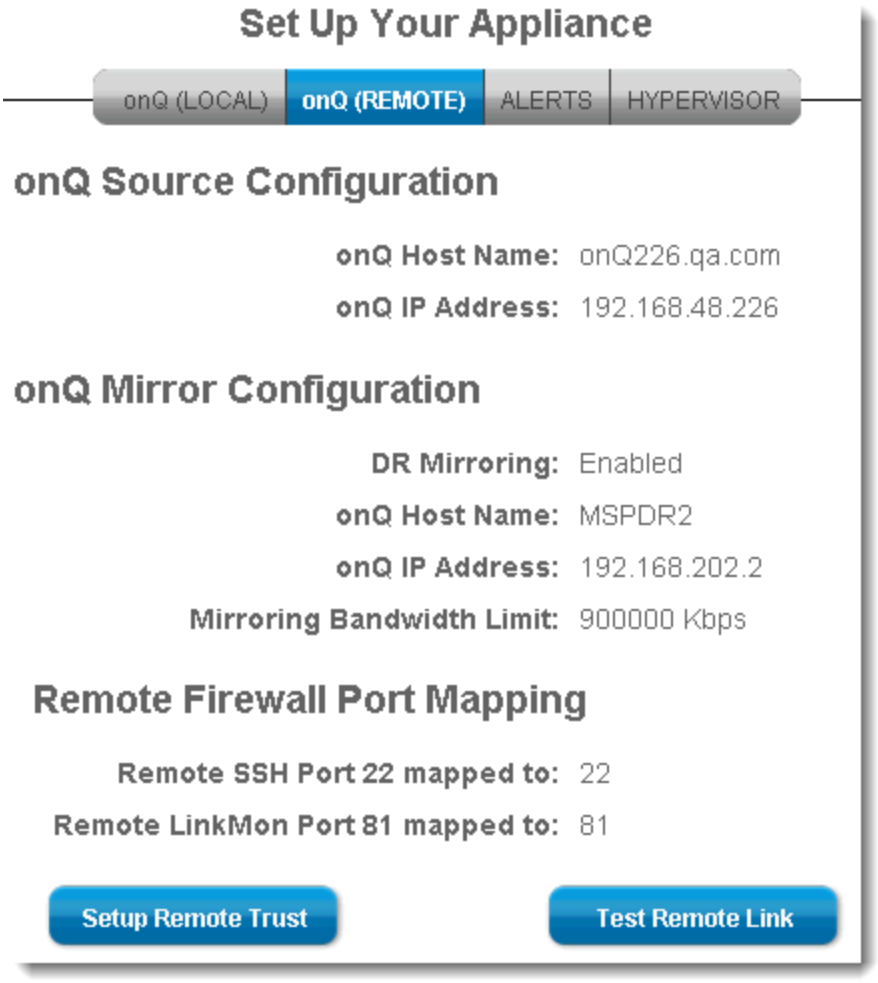
1. Log on to the HA Appliance's onQ Portal.
2. Go to **APPLIANCE CONFIG** tab > **onQ (REMOTE)** page.

   Simply click on the **Set Up Remote Trust** button.

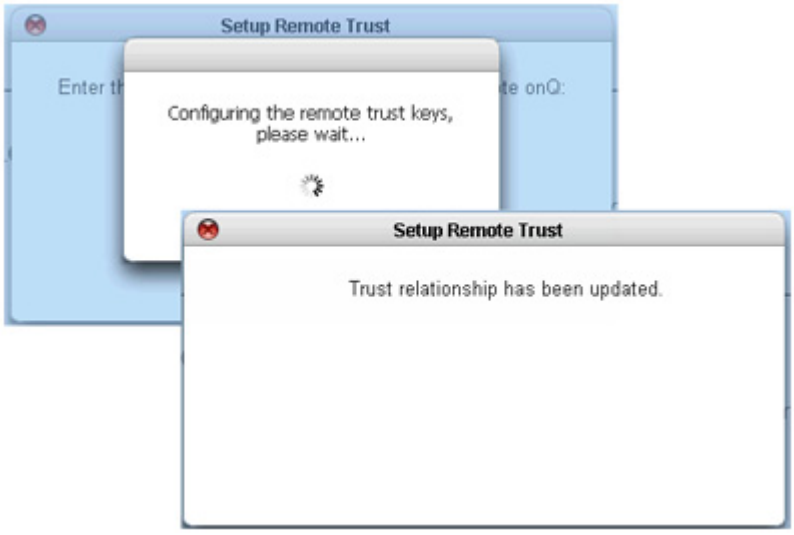
Example of HA Appliance:



Example of DR Appliance:



3. Type the VARAdmin password for the *other* onQ Appliance, then **Update Trust**.

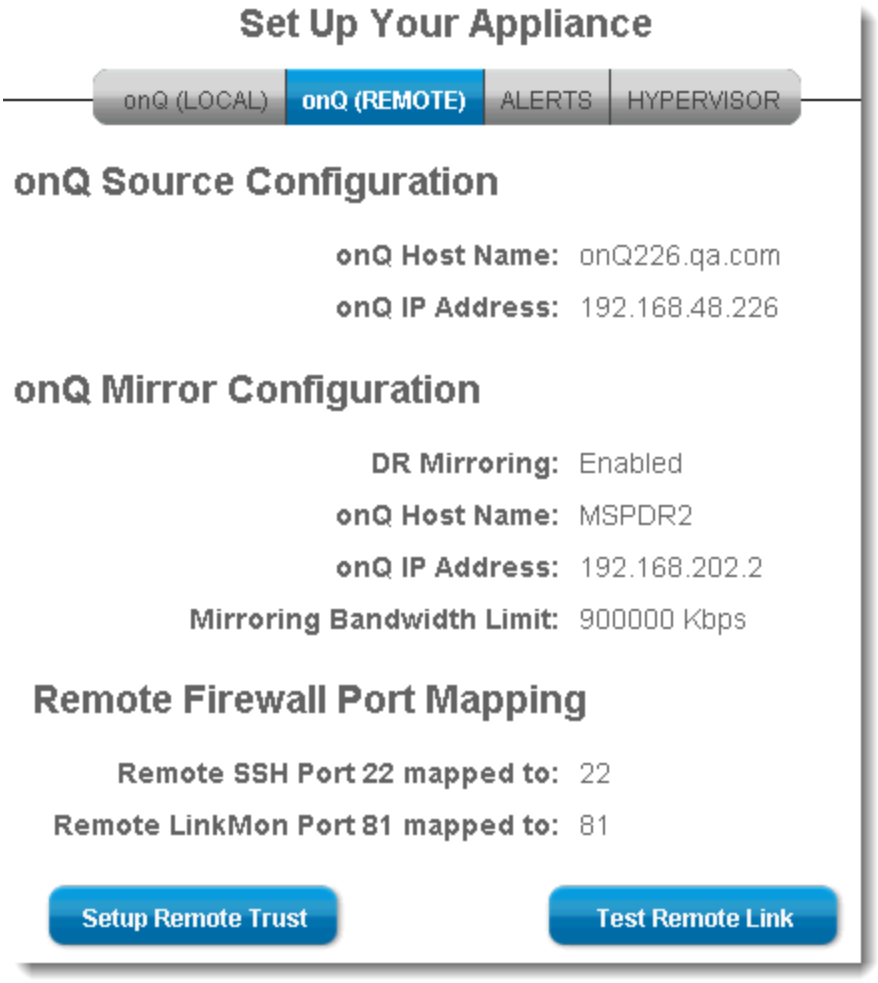


4. Test the link.

To set up the trust relationship between DR Appliance and DR Mirror:

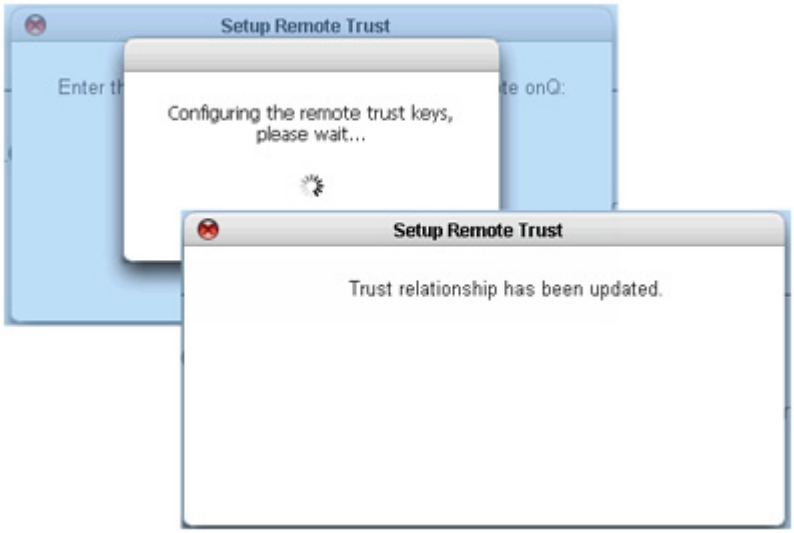
1. Log on to the DR Appliance's onQ Portal.
2. Go to **APPLIANCE CONFIG** tab > **onQ (REMOTE)** page.

Example of DR Appliance:



3. Simply click on the **Set Up Remote Trust** button. These buttons are inactive if you have not yet enabled DR mirroring.

4. Type the VARAdmin password for the *other* onQ Appliance, then **Update Trust**.



5. Test the link.

10.3 Create Custom Networks for RNs

If you currently segment your traffic using VLANs (virtual networks), you will want to replicate this configuration on your onQ. For example, your traffic might be segmented by internal and external (DMZ) traffic, by department, or iSCSI vs. SAN.

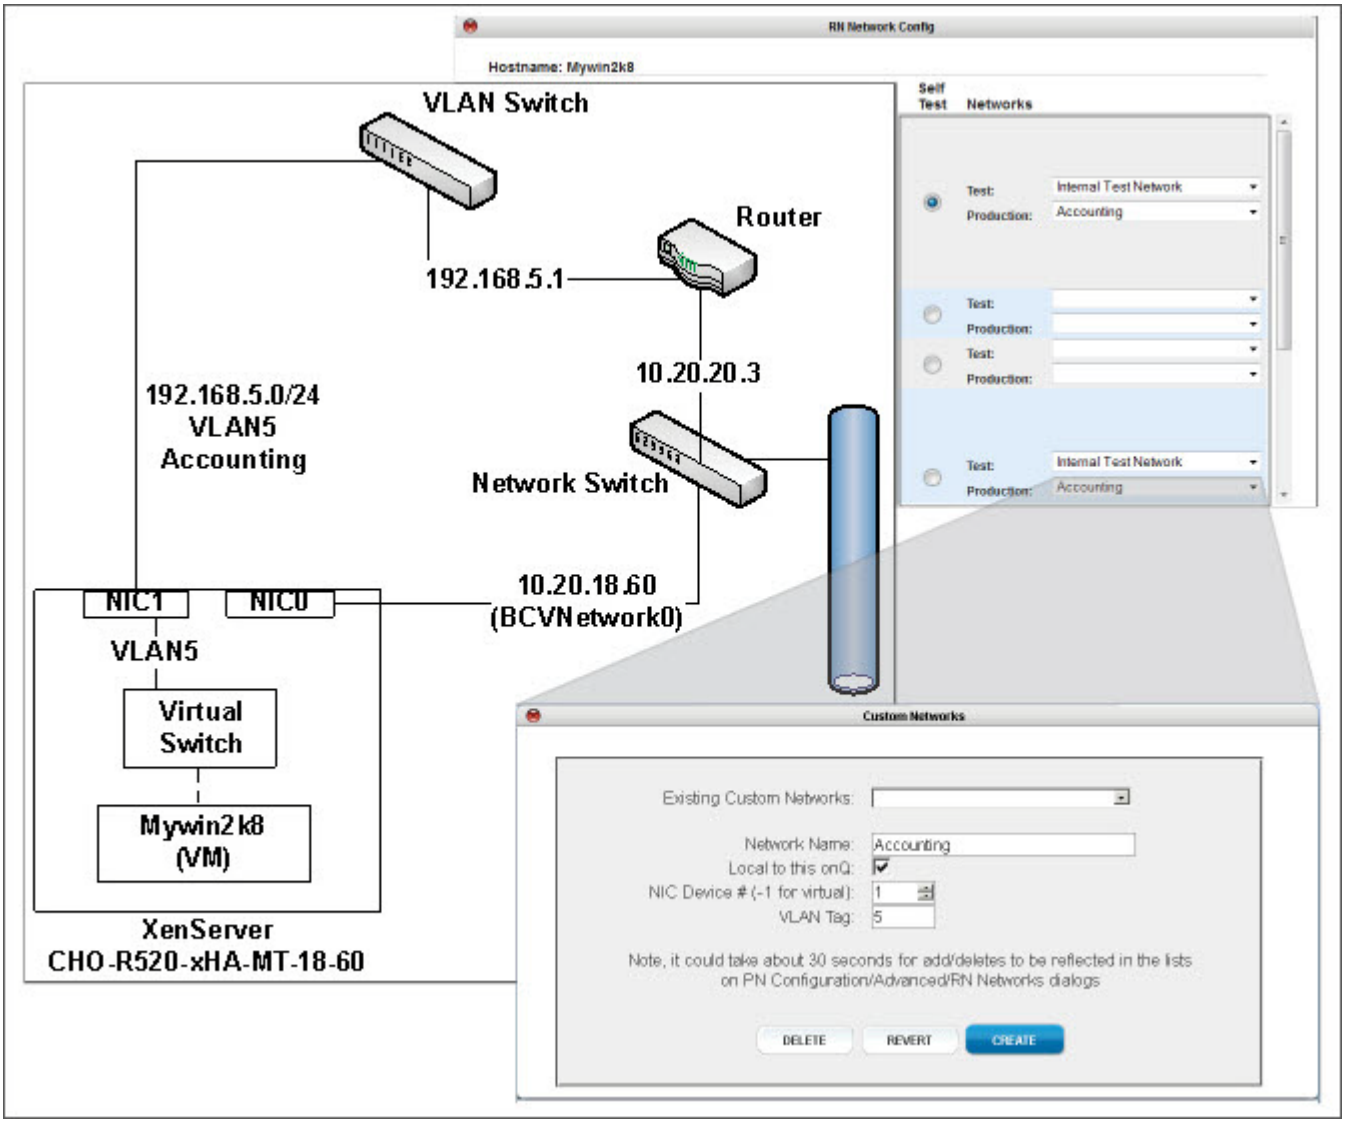
Quorum recommends that you set up these custom networks using the onQ Portal. Do not set up these custom networks on the onQ hypervisor directly. The onQ Portal provides all the bells and whistles to help you emulate your network easily and safely.

To create virtual networks:

1. Identify your broadcast domains (aka VLAN IDs or VLAN tag), the ports that they comprise, and the PNs that reside in each network.
2. Log on to the onQ Appliance's onQ Portal. You must log on as `varadmin` user. When you create a network on the HA, it does not automatically appear on the DR. You must make such changes on both onQ Appliances.

3. Click the **APPLIANCE CONFIG** tab > **HYPERVISOR** page.
4. Click the **CUSTOM NETS** button. The Modify Custom Networks dialog appears.
5. Specify the custom network parameters, click **CREATE**, then **SAVE**.
   - **Existing Custom Networks**. If you've previously defined a custom network using the onQ Portal, it appears in this list; however, if you previously configured custom networks on the onQ hypervisor directly, the onQ Portal does not present these networks.
   - **Network Name**. Specify a descriptive name that helps you identify this VLAN (or the type of traffic on this virtual network). For example, `DMZ Traffic` or `Accounting Traffic`. After you define your custom network, this network name appears in the **Existing Custom Networks** drop-down list.
   - **Local to this onQ**. Checking this check box means that the network is only available to RNs of this onQ instance. Clearing the check box allows the network that you're defining to be available to the RNs that belong to all onQ instances on the onQ Appliance.
   - **NIC Device # (-1 for virtual)**. Specify either `-1` for virtual NIC or `1-4094` for a physical NIC. You cannot leave this field blank.
   - **VLAN Tag**. Specify the unique identifier (aka VLAN ID or broadcast domain) for the virtual network. An identifier is inserted into a packet header in order to identify which VLAN the packet belongs to. Switches use the VLAN Tag to determine which port(s) or interface(s) to send a broadcast packet to. This tagging allows you to create multiple VLANs on the same physical interface.

Example: *Separate VLAN for Accounting Traffic*



6. Now that you've created the networks, assign your RNs to those networks as outlined in Assign RNs to Networks.

10.4 Assign RNs to Networks

If you have Windows PNs that are on different subnets, you'll want to configure multiple networks so that those RNs have a pathway outside the onQ Appliance. (Linux PNs are not yet supported.) Additionally, you can create complex test networks to test your RNs.

Quorum recommends that you set up multiple networks using the onQ Portal. Do not set up these networks on the onQ hypervisor directly. The onQ Portal provides all the bells and whistles to help you emulate your network safely and easily.

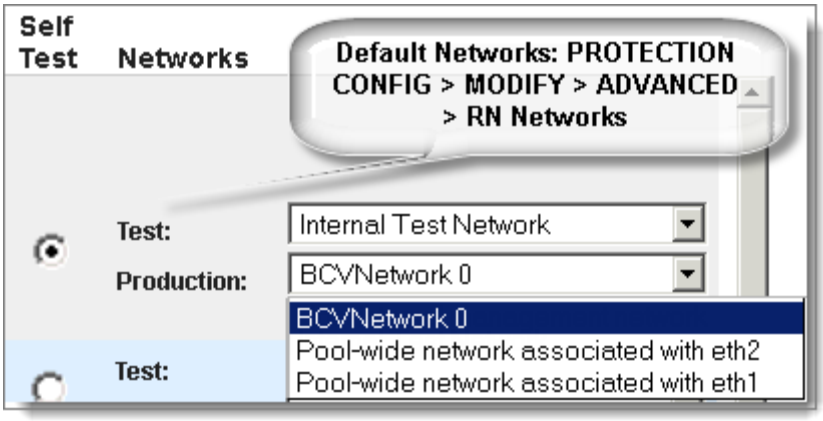
Your supported Windows RNs (see Platform Support) can run in either production mode or test mode on multiple networks using virtual NICs. An RN's IP address changes during an RN build to accommodate these different networks.

When an RN starts in production mode or test mode, it will connect to the pre-configured network. If you do not configure for multiple networks, the RN connects to the `BCV Network 0` (default) or `Internal Test Network` respectively; if you have PNs on different subnets, they will not have a pathway outside the onQ Appliance.

Your RNs can also be assigned to virtual networks, but you must set up these virtual networks in advance as outlined in Create Custom Networks for RNs.

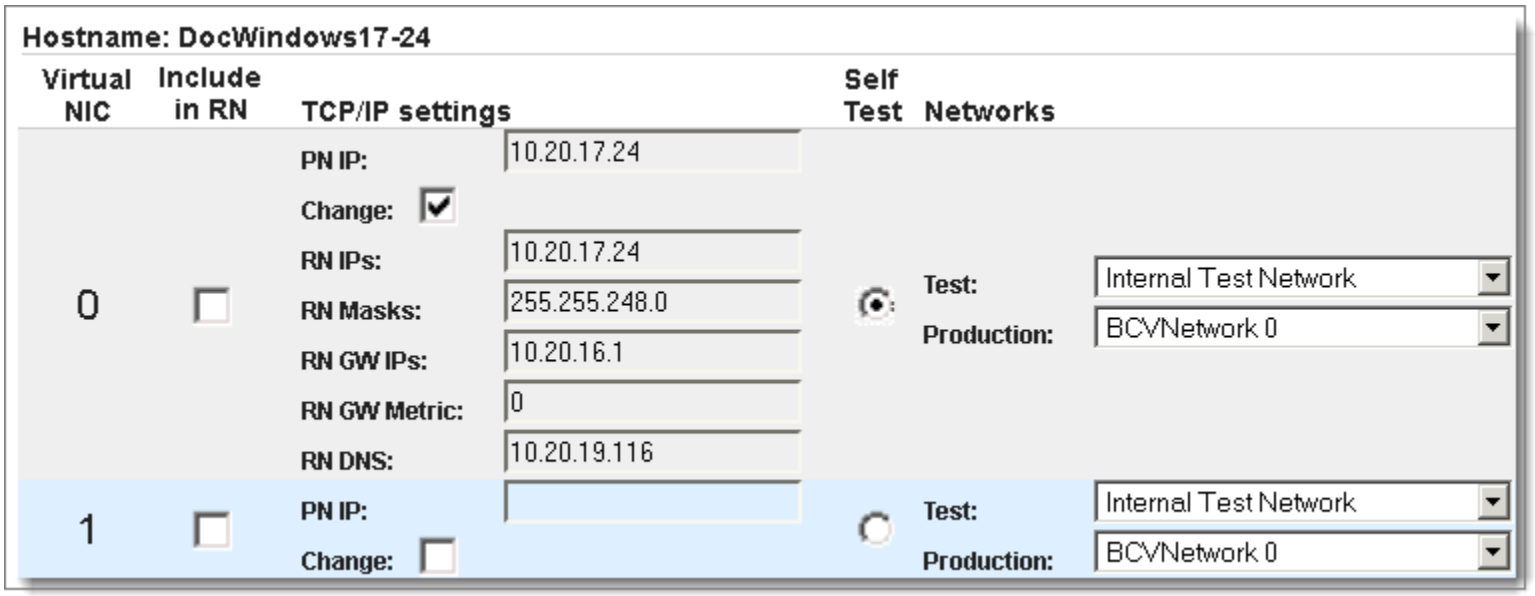
If you want to attach an RN directly to a physical NIC, you do not need to create a custom network. Simply choose either the existing `BCVNetwork 0` virtual network or one of the `Pool-wide` virtual networks that is associated

with `ethX`:



To assign RNs to networks:

1. Log on to the onQ Appliance's onQ Portal.
2. Click the **PROTECTION CONFIG** tab.
3. Select the node, then **MODIFY** button > **ADVANCED** button.
4. In the **RN Networks** field, click **Default** button.

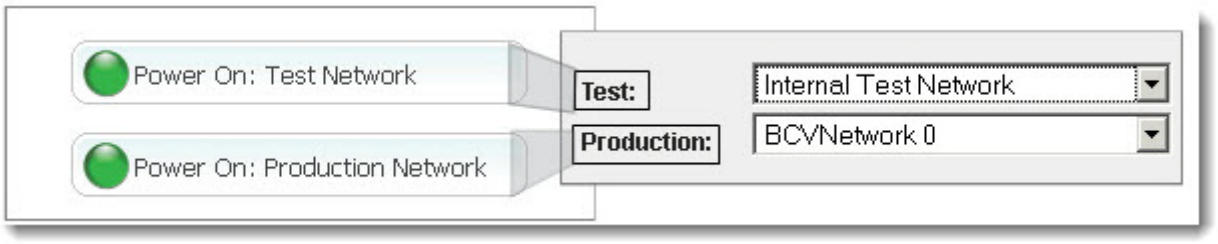


5. Configure the RN's network(s). The RN defaults to the PN's networking information, then **SAVE**.
   - **Virtual NIC**. Define your virtual NICs. The hypervisor supports up

to 7 virtual NICs (NIC 0-6). NIC 0 defaults to the PN's pre-configured TCP/IP settings. Use the **Include in RN** check box to make them available to the RN.

- **Include in RN**. Select the check box for each Virtual NIC that you want to make available to the RN.
- **TCP/IP settings**. Specify the networking information for each Virtual NIC that you want to make available to the RN. You can set different IPs for the HA and DR, enabling RNs to boot up with different IPs per location. RN Gateway metric is a 32-bit integer; the chosen route will be the gateway with the lowest metric.
- **Self Test**. Select the Self Test radio button that corresponds to Virtual NIC and TCP/IP settings that you want the onQ to use for all Self Tests.
- **Networks**. Select the test network and production network that you want onQ to use when running an RN in either mode. Networks include:

Network Name	Description
Test Networks	
Internal Test Network	Virtual adaptor for running RN in test mode and for running RN self-tests.
Production Networks	
Internal Test Network	Virtual adaptor for running RN in test mode and for running RN self-tests. Do not use as a production network.
BCVNetwork 0	NIC0
Pool-wide network associated with eth2	NIC1 on the onQ Appliance. Available to all VMs on that onQ Appliance.
Pool-wide network associated with eth1	NIC2 on the onQ Appliance. Available to all VMs on that onQ Appliance.


Power On: Test Network
Power On: Production Network
Test:
Production:
Internal Test Network
BCVNetwork 0

10.5 Modify TCP Ports for HA-DR Communications

You might need to modify the TCP port numbers used for communications between the local and remote onQ Appliances if inter-Appliance communications traverse a firewall that does network address translation (NAT) on the onQ Manager IP addresses.

- **Port 22** (ssh) — the default port for securely transmitting data. It's highly unlikely that you'd ever need to change this port.
- **Port 81** — the default port that each onQ Appliance uses to verify that the other is "alive"; no data is transmitted via this port.

To modify the TCP port numbers:

1. Stop protection.
2. Go to **APPLIANCE CONFIG** tab > **onQ (REMOTE)** page.
3. Click **MODIFY**.

   The **Modify Remote onQ Setup** page appears.
4. In the **Remote Firewall Port Mapping** section, make your changes, then **SAVE**.
5. Test the link.
6. Restart protection.

Related Topics

Network and Firewall Requirements

10.6 Synchronize system time

The HA Appliance uses an NTP Server for its system time (for more information, go to Configure Appliance's hypervisor settings).

The HA Appliance and the protected nodes must not be out of sync by more than 180 seconds during installation or 5 minutes after installation. If they are out of sync, a secure connection cannot be established.

Symptoms of an unsecure connection are outlined in (Agent-based PNs) Connection Problems.

To verify the system time on the onQ Appliance:

1. Log on to either HA Appliance's onQ Portal or the DR Appliance's onQ Portal.
2. Go to **DASHBOARD** tab > **onQ STATUS** page. The time appears as a header on the page.

To synchronize the clocks:

If the clock on the onQ Appliance is not synchronized with the protected nodes, you have a few choices:

- **(Best)** Enable NTP on the protected nodes or the Domain Controller.
- **(Better)** If your protected node is acquiring its time from your Domain Controller, update the time on that server.
- **(Not Recommended)** Update the system time on each individual protected node to be in sync with the onQ Appliance. This process is tedious, but it works.

10.7 Secure connection to PNs

The network connection between the HA Appliance and the PN is generally over a trusted LAN. In some cases, though, it might be desirable to provide strong authentication between the HA Appliance and the PN.

The onQ Appliance uses an X.509 certificate for security. Your onQ Appliance ships with a unique certificate already installed. onQ automatically copies over this certificate to the PNs when you enroll protected nodes and again when you upgrade them, establishing a secure connection between the onQ Appliance and the PNs. The onQ portal's Connection Status column uses a green padlock icon to indicate a secure connection as outlined in Monitor protected nodes. The certificate has a long expiration.

You can regenerate and/or reinstall a new certificate at any time. There are a couple of scenarios where you might want to do so:

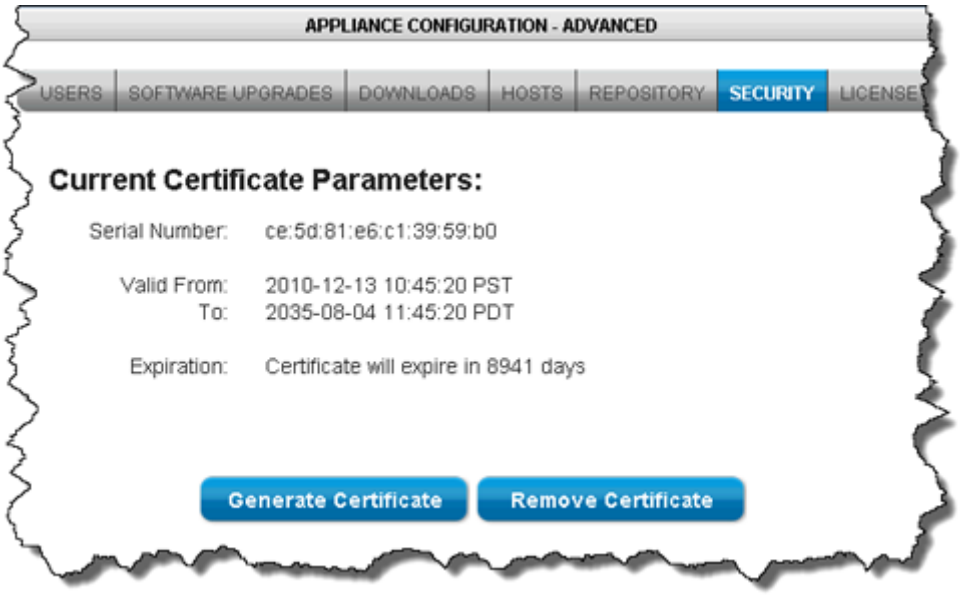
- Your HA Appliance's certificate is corrupt; therefore, so too is the certificate on your PNs.

- You experienced a disaster scenario that resulted in you having to change the role of the DR Appliance to `HA`. In order for the RNs to establish a connection to that DR Appliance, you need to install the DR Appliance's certificate on those RNs.

To fix a corrupt certificate:

Fixing a corrupt certificate on requires that you regenerate and reinstall a new certificate. This procedure assumes that the HA Appliance's certificate is corrupt, but you can extrapolate.

1. Log on to the HA Appliance's onQ Portal.
2. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **SECURITY** page.
3. Click **Remove Certificate**, then **Yes** to remove the existing certificate on the onQ Appliance. An example certificate is as follows:



4. Click **Generate Certficate**.
5. For each PN, re-install the HA Appliance certificate on the PNs by launching the HA Appliance's onQ Portal from each PN and re-enrolling the protected nodes using the **Protect Me** button.

To establish a secure connection between RNs and a DR Appliance:

Let's assume that you experienced a disaster scenario that resulted in you having to change the role of the DR Appliance to `HA`. (If you're performing a

failback, extrapolate). In order for the RNs to establish a connection to that DR Appliance, you need to install the DR Appliance's certificate on those RNs because the certificate that was originally installed on the PNs is specific to the HA Appliance, which has a different hostname than the DR Appliance, the current acting HA.

To do so, simply launch the DR Appliance's onQ Portal from the PNs, then re-enroll the protected nodes using the **Protect Me** button.

To manually install the certificate:

onQ automatically installs the certificate that you generate when you enroll protected nodes and again when you upgrade them. However, you can manually install the certificate, if you desire.

1. Log on to the HA Appliance's onQ Portal.
2. RDP to the protected node.
3. From that PN, launch the HA Appliance's onQ Portal.
4. Go to **PROTECTION CONFIG** tab, then click on the **Update PN Security** button.
5. Click **Yes** to download the file to the protected node. Your browser saves the file to its default location.
6. Move the *onQApplianceName*`.cert` file to **C:\Program Files\Quorum\QuorumDCRM-NODE\security**.
7. Restart the onQ Service.
8. Restart protection.

11

Disaster Recovery and DR Preparedness

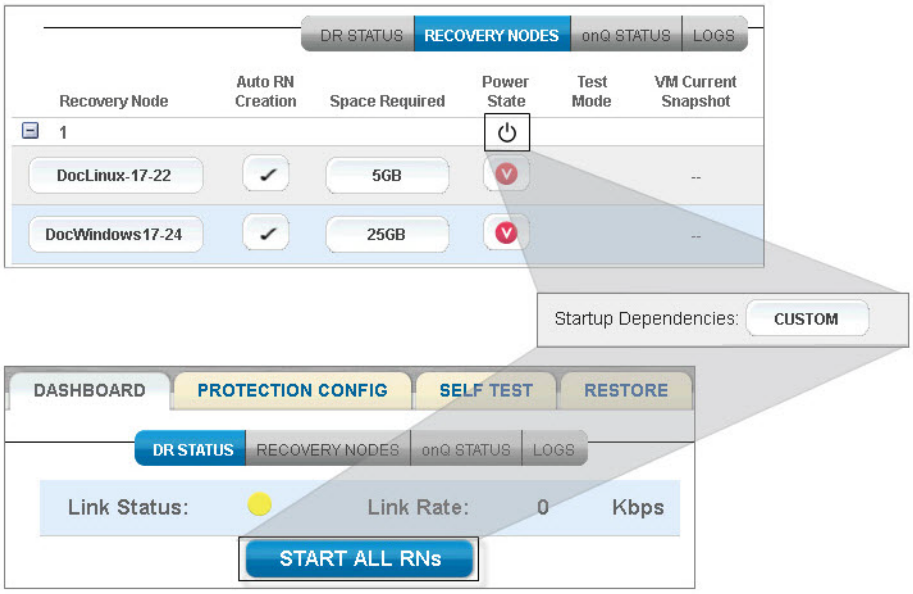
- Configure startup dependencies
- (Workflow) Fail over HA Appliance to DR Appliance
- (Workflow) Fail back DR to HA Appliance
- (Workflow) Fail over a PN to an RN
- (Workflow) Fail back an RN to a PN
- Change Appliance's role
- Start recovery nodes on HA Appliance
- Start recovery nodes on DR Appliance or DR Mirror
- Stop recovery nodes
- Back up and restore Oracle 11g database on Linux
- Back up and restore Oracle 10g+ database on Windows
- Cancel or override RN startup process
- Launch recovery node's console
- Test Mode
- Configure automatic testing of RNs
- Suspend and resume RN tests
- Restore PNs Using QUARK

11.1 Configure startup dependencies

If an RN depends on one or more RNs, or server outside your HA Appliance's control, (for example, your Exchange Server depends on a Domain Controller), configure these dependencies in advance. You can start all your recovery nodes individually, taking into account the necessary boot order, but onQ automates all this work for you.

onQ will start a given RN after the RNs on which it depends are up and running. Startup dependencies do not apply to RNs in test mode or automated self-test. Startup dependencies apply in production mode.

onQ uses startup dependencies when a group of RNs is started on the HA or a group, or all, RNs are started on the DR Appliance or DR Mirror. For more information, go to "Start recovery nodes on HA Appliance" on page 277 and "Start recovery nodes on DR Appliance or DR Mirror" on page 285.



To configure an RN's startup dependencies:

There is no limit to the number of dependencies that you can create for each RN. All dependencies in the list are executed in parallel.

All dependencies are configured on the HA Appliance. After you save that configuration, onQ transfers that information to the DR Appliances and DR Mirrors.

1. Log on to the HA Appliance's onQ Portal.
2. Create a group to contain each dependency. For example, create a **`Domain Controller & Mail Server`** group and assign both the Exchange Server and the Domain Controller to that group.
   a. Click the **PROTECTION CONFIG** tab.
   b. Select the node, then **MODIFY** button.
   c. In the **Group Name** field type the name of the group, then **SAVE**.



3. Specify the startup dependencies for the RN:
   a. Click the **PROTECTION CONFIG** tab.
   b. Select the node, then **MODIFY** button > **ADVANCED** button.
   c. In the Startup Dependencies field, click the **NONE** button. If this isn't your first time setting up a dependency, the **CUSTOM** button appears instead.
   d. Click the plus button (**+**). A default entry appears in the Startup Dependencies List.
   e. Define the **Dependent Server**. In the Dependent Server field, select the RN from the drop-down list or specify the IP address or hostname of the server that is not protected by (external to) this HA Appliance and on which the RN depends. If the dependent server is an external server, you must define it in the hosts file, unless you specify the IP address. Leave this field blank to specify a delay policy instead of a server dependency (see "**Delay**" in Step f).
   f. Specify the following information:

      Timeout – Indicate how long you want the onQ to wait for the **Dependent Server** to respond to a ping request. The default value is 5 minutes, but base the value on the typical boot time for your

server. Ensure that the firewall allows ping requests. If you specify a Dependent Server, you must also specify a timeout value.

Timeout Action – Indicate what you want onQ to do with the RN if the **Dependent Server** does not respond to the ping command within the timeout period. onQ can start the RN (`Start Anyway`) without honoring the dependent server, or you can instruct onQ to refrain from starting the RN (`Abort Start`), regardless of whether all other dependencies are met, because the dependent server is critical to the RN's availability.

Delay – Specify a delay policy along with or instead of a server dependency.

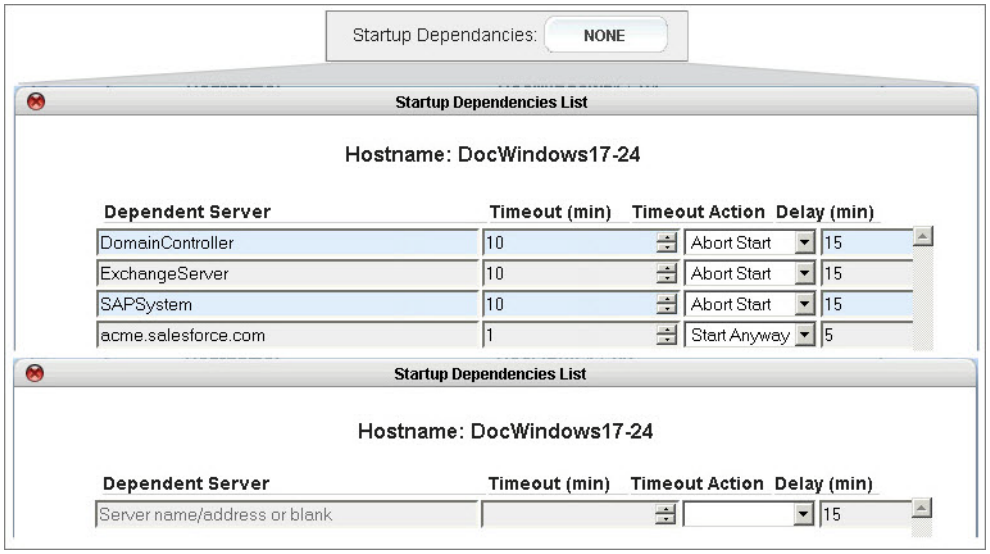
– **Delay policy *in addition* to a server dependency**. Indicate how long you want onQ to delay the RN startup after onQ detects that the dependent server is up. The default value is 1 minute, but base the value that you specify on the typical time needed for your dependent server's *applications* to respond to user requests.

– **Delay policy *instead of* a server dependency**. If a dependent RN is not configured to respond to the ping command on which **Timeout** depends, or you know that a dependent server will take a long time to boot, do not specify a server dependency; instead, create a delay policy: leave the **Dependent Server** field blank and specify a **Delay** time that accounts for the typical time needed for your dependent server's *applications* to respond to user requests. In this case, the delay option ignores values for **Timeout Action** and **Timeout**. You can only specify one such delay policy; if you try to specify two policies, the onQ Portal returns an error: `only one row may have a blank server entry for a delay only specification`.

Example:

In the following example, all the internal dependent servers typically require 5 minutes to respond to ping requests, but this example provides an additional 5 minutes just in case. In most cases, the dependent server's applications are available within 15 minutes; therefore, in this example, the applications will have more than enough time (10-min timeout plus a 15-min delay) to become available. All dependent servers are critical to the RN's availability, with the exception of the "optional" Salesforce application.

Alternatively, a simple 15-minute (or 20) delay policy can encompass all the dependencies in the list.



4. Click **SAVE**.
   - To revert to the last saved values, click **REVERT**.
   - To clear all data and start over, click **CLEAR**. Default values repopulate the fields provided.

   As the RN starts and if that RN has dependencies configured, the onQ Portal displays messages as the dependencies are met.

Related Topics

Cancel or override RN startup process

Start recovery nodes on HA Appliance

11.2 (Workflow) Fail over HA Appliance to DR Appliance

Use this procedure in the event of a site disaster at the HA site and if you expect the site to be down for a significant period of time, or if you've experienced an onQ Appliance failure and are waiting for a replacement.

If you have a cloud DR Appliance (see About Hybrid Cloud), Quorum can ship you the *cloud* DR Appliance, replacing it with a new DR instance. This *onsite* DR Appliance assumes the HA role.

Only an HA Appliance (local) can back up recovery nodes. Because your HA Appliance is unavailable, the DR Appliance (remote) must take on this role; otherwise, your recovery nodes, which are now acting as the protected nodes, will not be backed up.

Therefore, a failover involves the DR Appliance taking on the HA role and running the recovery nodes.

As soon as your local site is available, you must fail back.

To fail over to remote onQ Appliance:

1. Ensure that the local onQ Appliance is down (powered off) or that protection is off.

   When an onQ Appliance boots, protection is automatically off. Protection on an HA Appliance instructs the HA Appliance to perform backups. Given that the HA Appliance (local) still has its role set to `HA`, protection on that HA Appliance must be off until the failback completes.

2. Log on to the DR Appliance's onQ Portal.
3. Start all the recovery nodes on the DR Appliance (remote onQ Appliance). The boot order must take into account any interdependencies.
4. If applicable, restore the Oracle database:
   - Back up and restore Oracle 11g database on Linux
   - Back up and restore Oracle 10g+ database on Windows
5. Add a host entry for each protected node, if you haven't already done so.

6. Change the DR Appliance's role to HA.

   The recovery nodes on the remote onQ Appliance become the protected nodes; the remote onQ Appliance backs up these recovery nodes.

7. Establish a secure connection between the RNs and the DR Appliance. Go to Secure connection to PNs.

Related Topics

(Workflow) Fail back DR to HA Appliance

11.3 (Workflow) Fail back DR to HA Appliance

Use this procedure if you performed the failover steps outlined in (Workflow) Fail over HA Appliance to DR Appliance.

To fail back to local onQ Appliance:

1. Determine how you'd like to synchronize the local onQ Appliance's repository with the remote onQ Appliance's repository, which has the most recent data.

   If there is too much data consider an export/import instead of a WAN synchronization.

2. Do one of the following based on the decision you made in Step 1.
   - If you want to synchronize the Repositories using WAN synchronization, ensure that the trust relationship is set up on the *remote* onQ Appliance.
   - If you want to synchronize the Repositories using export/import, remove the trust relationship on the *remote* onQ Appliance by setting an empty password.
3. Schedule downtime based on the time needed to sync or export/import the repository data.
4. During the scheduled downtime, log on to the *remote* onQ Appliance's onQ Portal and do the following:
   a. Stop all the recovery nodes, which are acting as protected nodes.
   b. Stop protection.
5. Power on the *local* onQ Appliance, if it isn't already.
6. Change the local onQ Appliance's role to HA, if it isn't already.
7. Do one of the following:
   - If you chose WAN synchronization in Step 2, that process takes place as soon as the local onQ Appliance boots up. At this time, also ensure that the local onQ Appliance's role is HA.
   - If you chose export/import in Step 2, ensure that the local onQ Appliance's role is HA, then perform that export/import.

   As the synchronization or export/import is in progress, proceed to Step 8.

8. Log on to the *local* onQ Appliance's onQ Portal, and do the following to address failures to the protected nodes, if applicable.
   a. Disable the protected nodes, so that the onQ Appliance does not perform backups during a restore.
   b. Start protection. A restore requires protection.
   c. Restore the protected nodes using one of the following methods:
      - bare metal restore (BMR) or incremental/reversion restore
      - file level restore (FLR)
      - Windows Share Restore (WSR)
9. On the *local* onQ Appliance, bring up all the protected nodes. The boot order must take into account any interdependencies.
10. Change the remote onQ Appliance's role to DR.
11. Restart protection.
12. Establish a secure connection between the PNs and the HA Appliance. Go to Secure connection to PNs.

Related Topics

(Workflow) Fail over HA Appliance to DR Appliance

11.4 (Workflow) Fail over a PN to an RN

When you start an RN, whether agent-based or agent-less, you are indirectly shutting down the PN and PN proxy (VM) respectively and disabling backups for that PN.

To fail over a PN to an RN *on the HA Appliance*:

1. Log on to the HA Appliance's onQ Portal.
2. Start the recovery node or, if the PN has a Build-on-Demand RN creation policy, create the on-demand RN.

   If the RN is agent-less, the Backup Mode switches from **`Agentless`** to **`Recovery Node`**.
3. If applicable, restore the Oracle database:
   - Back up and restore Oracle 11g database on Linux
   - Back up and restore Oracle 10g+ database on Windows
4. Fix the PN's server.

When you're ready to fail back to the PN, go to To fail back an agent-based RN to a PN on the HA Appliance:

To fail over a *agent-less* PN to an RN *on the DR Appliance*:

If you start an agent-less PN's RN in production mode on the DR Appliance, you must perform the following steps so that the HA Appliance can back up the RN in production mode on the DR Appliance. This procedure enables onQ to back up the RN on the HA Appliance as a standard agent-based PN.

1. Start the RN in production mode on the DR Appliance.
2. RDP to the RN, then, from a browser window, launch the HA Appliance's onQ Portal.
3. Click the **Protect Me** button to install the onQ Monitor service on the RN.

   In the **Modify a Protected Node** page, switch the RN's Backup Mode from **`Agentless`** to **`Recovery Node`**, then **SAVE** the configuration.

When you're ready to fail back to the PN, go to To fail back an agent-less RN to a PN on the DR Appliance:

Related Topics

(Workflow) Fail back an RN to a PN

11.5 (Workflow) Fail back an RN to a PN

This workflow assumes that your enrolled PN went down for some reason, and that you failed over the PN to an RN as outlined in (Workflow) Fail over a PN to an RN. Now that you've fixed your original PN, use the failback procedure that pertains to your PN's enrollment method to restore the original PN:

To fail back an *agent-based* RN to a PN *on the HA Appliance*:

1. Log on to the onQ Appliance's onQ Portal.
2. Change the onQ Proxy address of the PN, if you specified an RN network configuration different than the original PN as part of multi-network support.
3. Schedule some downtime. Prevent any further changes to the RN.
4. Perform a backup of the RN in production mode.
5. Shut down the RN in production mode.
6. Perform an incremental restore of the PN, choosing the latest snapshot that you initiated in Step 4.

   onQ resumes backups of this agent-based PN.

To fail back an *agent-less* RN to a PN *on the HA Appliance*:

1. Log on to the onQ Appliance's onQ Portal.
2. Manually install the node software package on the RN.
3. Change the onQ Proxy address of the PN, if you specified an RN network configuration different than the original PN as part of multi-network support.
4. Schedule some downtime.
5. Perform a backup of the RN in production mode.
6. Shut down the RN in production mode.

   The Backup Mode switches from **`Recovery Node`** to **`Agentless`**.
7. Boot the original PN using QUARK.
8. Perform an incremental restore, choosing the latest snapshot that you initiated in Step 5.

onQ resumes backups of this agent-less PN.

To fail back an *agent-less* RN to a PN *on the DR Appliance*:

1. Launch the DR Appliance's onQ Portal.
2. Shut down the RN in production mode on the DR Appliance.
3. Restart the original PN in production mode on the DR Appliance.
4. Re-enable backups of the PN:
   a. Go to **APPLIANCE CONFIG** tab > **Modify a Protected Node** page.
   b. Select the PN, then **MODIFY**.
   c. Switch the PN's Backup Mode from `Recovery Node` to `Agentless`, then **SAVE** the configuration.

11.6 Change Appliance's role

An onQ Appliance can have one of two roles: **HA** (for high availability) or **DR** (for disaster recovery). For a new onQ Appliance, the default is HA.

You'll need to change the role on the on-premise DR Appliance to HA if you experience a failure as outlined in (Workflow) Fail over HA Appliance to DR Appliance.

If, for example, you set the onQ Appliance's role to HA, then **onQ role: HA** displays at the top of each UI page on that onQ Appliance. Until you change the other onQ Appliance's role to DR, both Appliances indicate HA.

To change an onQ Appliance's role:

1. Log on to the onQ Appliance's onQ Portal.
2. Go to **APPLIANCE CONFIG** tab > **onQ (LOCAL)** page.

   You will see a variety of settings for this onQ Appliance.
3. Click **MODIFY**.
4. In the **onQ Role** drop-down list, choose the role you want to assign.

   A window appears, asking you to reboot the onQ Appliance.
5. Reboot onQ.

Related Topics

Secure connection to PNs

11.7 Start recovery nodes on HA Appliance

In the event of a disaster, you might need to start an individual recovery node or a group of recovery nodes on HA Appliance.

Use this procedure in conjunction with the appropriate failover or failback procedures.

- (Workflow) Fail over HA Appliance to DR Appliance
- (Workflow) Fail back DR to HA Appliance
- (Workflow) Fail over a PN to an RN
- (Workflow) Fail back an RN to a PN

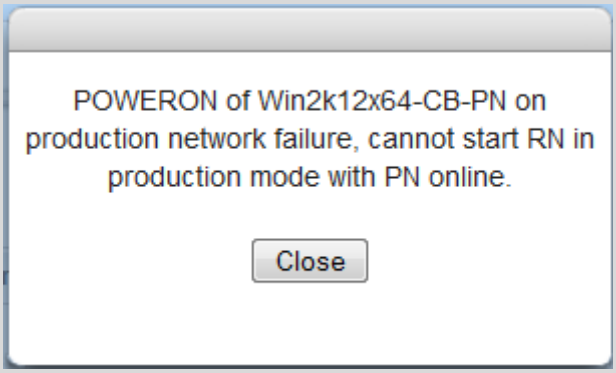
When you start an RN, whether agent-based or agent-less, you are indirectly shutting down the PN and PN proxy (VM) respectively and disabling backups for that PN.

On an HA Appliance there are two ways to start *individual* recovery nodes or a *group* of recovery nodes:

- **in production mode** — use this mode if your primary protected node failed and you want to immediately start the recovery nodes in your production environment. After you fix the PN, don't forget to stop the recovery node before bringing the PN online; otherwise, there will be network conflicts.

- **in** test mode — use this mode if you simply want to verify that the recovery node will run in the event of a disaster or to test a host running new server updates.

Note: So as to avoid network conflicts, if the corresponding PN is already running in production, the onQ Portal will not let you start the RN in production. In that case, the onQ Portal displays the following error message:



You may be prompted to install application-specific software licenses upon startup of recovery nodes.

If you want to start *all DR* recovery nodes, go to Start recovery nodes on DR Appliance or DR Mirror.

(From RN tab) To start a *group* of recovery nodes in *production mode* or *test mode*:

1. From the HA Appliance, go to **DASHBOARD** tab > **RECOVERY NODES** page.
2. Unlock the page.
3. For the group, click the **Power State** (⏻) button.
4. Select either **Power On: Test Network** or **Power On: Production Network**.

   When you start in production mode, the startup order will take into account the order you specified in "Configure startup dependencies" on page 264.

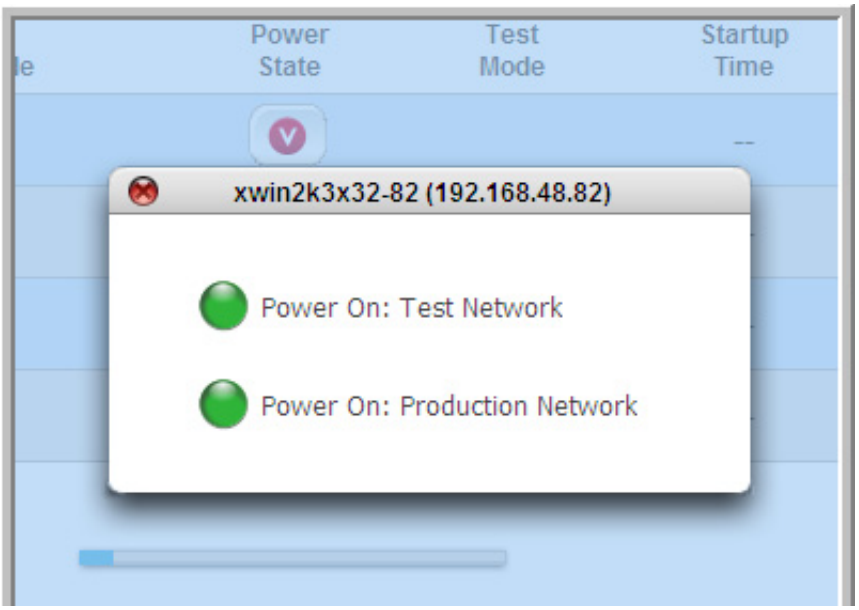
(From RN tab) To start an *individual* recovery node in *production mode*:

1. (Optional) Consider suspending RN tests to "free up" resources.
2. (Redhat 5.x RNs) Install kernel-xen RPM package.
3. Shut down the corresponding PN, if it's online.
4. From the HA Appliance, go to **DASHBOARD** tab > **RECOVERY NODES** page.
5. Unlock the page.
6. Click the button in the **Power State** column, then choose to enable or disable backups.

   A dialog box appears that offers you a choice to disable backups for the RN (aka Recovery PN).
   - If you're trying to recover from a failed PN, it's important that backups continue.
   - If you are testing an RN in production after applying server updates, you might want to disable backups.

Another dialog box appears that offers you a choice between the following two alternatives: **Power On: Test Network** or **Power On: Production Network**.

7. Choose the **Power On: Production Network** power option.



8. If you have an Oracle database, you must restore it now. Go to:
   - Back up and restore Oracle 11g database on Linux.
   - Back up and restore Oracle 10g+ database on Windows

(From PN tab) To start an *individual* recovery node in *production mode*:

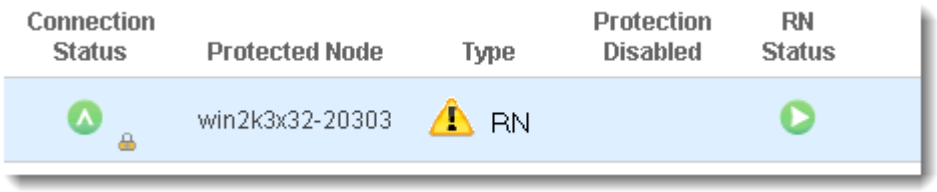
1. (Optional) Consider suspending RN tests to "free up" resources.
2. (Redhat 5.x RNs) Install kernel-xen RPM package.
3. Go to **DASHBOARD** tab > **PROTECTED NODES** page.
4. Unlock the page.
5. Click on the **Connection Status** button for the recovery node that you want to start.

   An information box appears that offers you this option: **Power ON the Recovery Node on the Production Network**.
6. Click **Continue**, to start the recovery node. If the recovery node has already started, this request does nothing.

   During the start up of the RN, the power indicator turns yellow with a green arrow:

   Soon after, an exclamation icon appears to the right of the protected node and the **Type** field changes to **RN**, indicating that the recovery node is now running.

7. if you have an Oracle database, you must restore it now. Go to:
   - Back up and restore Oracle 11g database on Linux.
   - Back up and restore Oracle 10g+ database on Windows

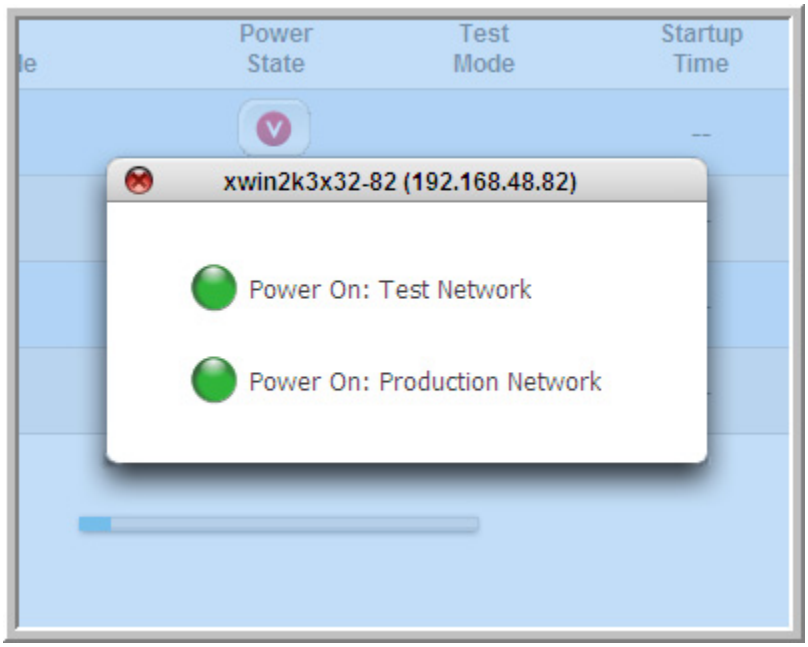
To start an *individual* recovery node in *test mode*:

1. Go to **DASHBOARD** tab > **RECOVERY NODES** page.
2. Unlock the page.

3. Click the button in the **Power State** column.

   A box appears that offers you a choice between the following two alternatives: **Power On: Test Network** or **Power On: Production Network**.

4. Click on the **Power On: Test Network** power option.



   A check mark appears in the **Test Mode** column of the appropriate status page, indicating that the recovery mode is running in test mode.

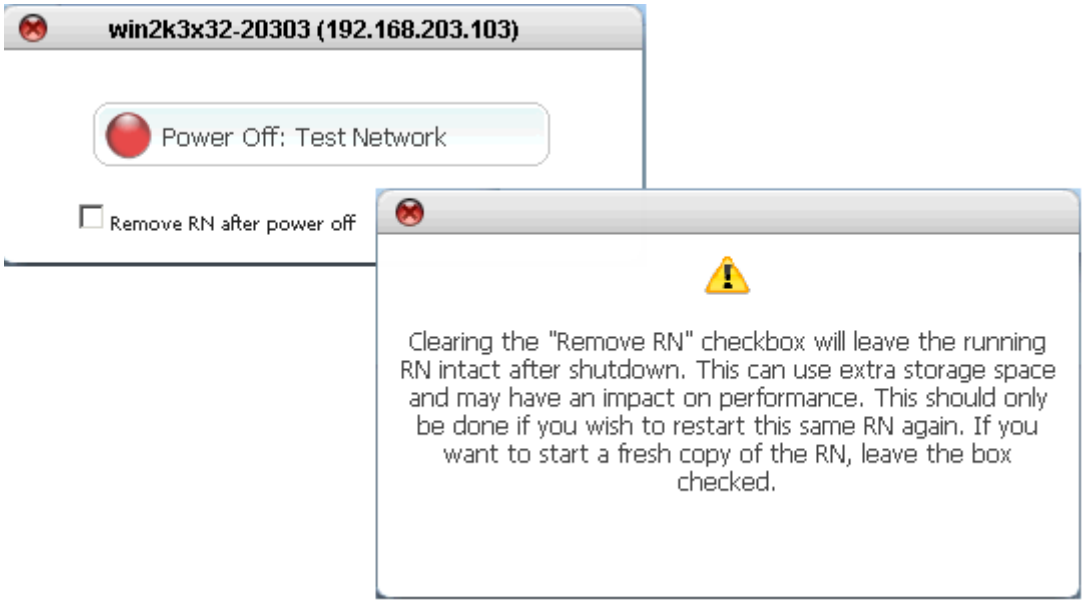
5. Launch the host's console.

6. Perform your testing.

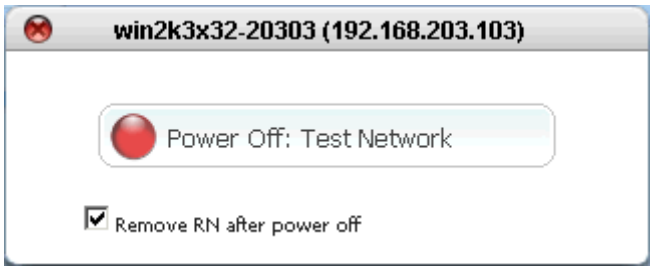
7. After you test the node, return it to production:

   a. Click the button in the **Power State** column.

   b. (Optional) Clear the **Remove RN after power off** check box if you want to retain this RN, keeping in mind that it is outdated.

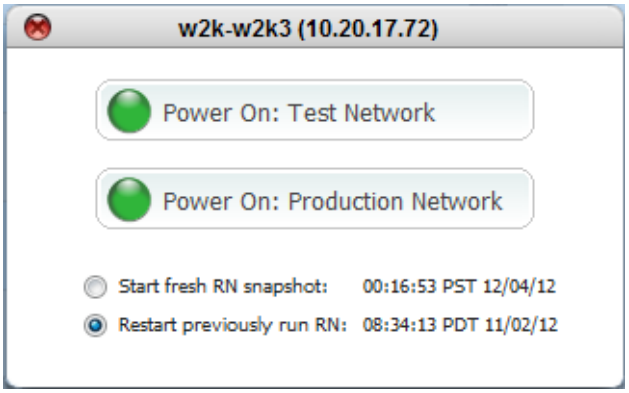


c. Click **Power Off: Test Network** option.



The icon in the **Power State** column changes color and orientation. There might be a delay of several seconds or minutes before the icon shows a change.

If you cleared the **Remove RN after power off** check box, then the following dialog will open the next time the RN is started.



8. Choose **Start fresh RN snapshot** or **Restart previously run RN** as appropriate.

Related Topics

Self-Test Alerts
Cancel or override RN startup process
Create on-demand RNs

11.8 Start recovery nodes on DR Appliance or DR Mirror

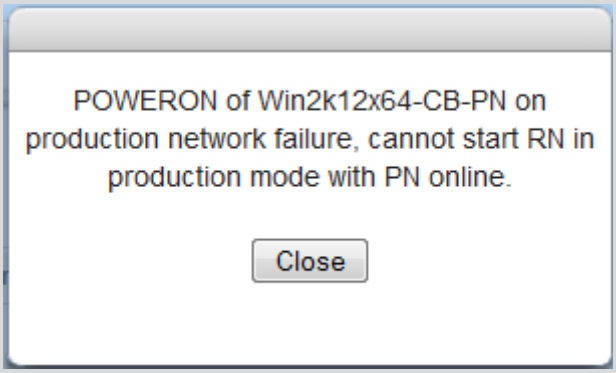
In the event of a disaster, you might need to start an individual recovery node, a group of recovery nodes, or all the recovery nodes on the DR Appliance.

Use this procedure in conjunction with the appropriate failover or failback procedures.

- (Workflow) Fail over HA Appliance to DR Appliance
- (Workflow) Fail back DR to HA Appliance
- (Workflow) Fail over a PN to an RN
- (Workflow) Fail back an RN to a PN

When you start an RN, whether agent-based or agent-less, you are indirectly shutting down the PN and PN proxy (VM) respectively and disabling backups for that PN.

Note: So as to avoid network conflicts, if the corresponding PN is already running in production, the onQ Portal will not let you start the RN in production. In that case, the onQ Portal displays the following error message:



It's possible to allocate resources to recovery nodes beyond what is physically available on the onQ Appliance. If you don't have enough memory to start all the recovery nodes, you have two choices:

- You can change the memory allocation of PNs so that the total is less than the memory on the onQ Appliance.

- You can start your recovery nodes individually. The boot order must take into account any interdependencies.

To start *all* recovery nodes in *production mode*:

1. (Optional) Consider suspending RN tests to "free up" resources.
2. (Redhat 5.x RNs) Install kernel-xen RPM package.
3. From the DR Appliance, go to **DASHBOARD** tab > **DR STATUS** page.
4. Unlock the page.
5. Click on the **START All RNs** button.



6. Click **Apply** to start all the DR recovery nodes.

   When you start in production mode, the startup order will take into account the order you specified in "Configure startup dependencies" on page 264.

To start a *group* of recovery nodes in *production mode* or *test mode*:

1. From the DR Appliance, go to **DASHBOARD** tab > **RECOVERY NODES** page.
2. Unlock the page.
3. For the group, click the **Power State** (⏻) button.
4. Select either **Power On: Test Network** or **Power On: Production Network**.

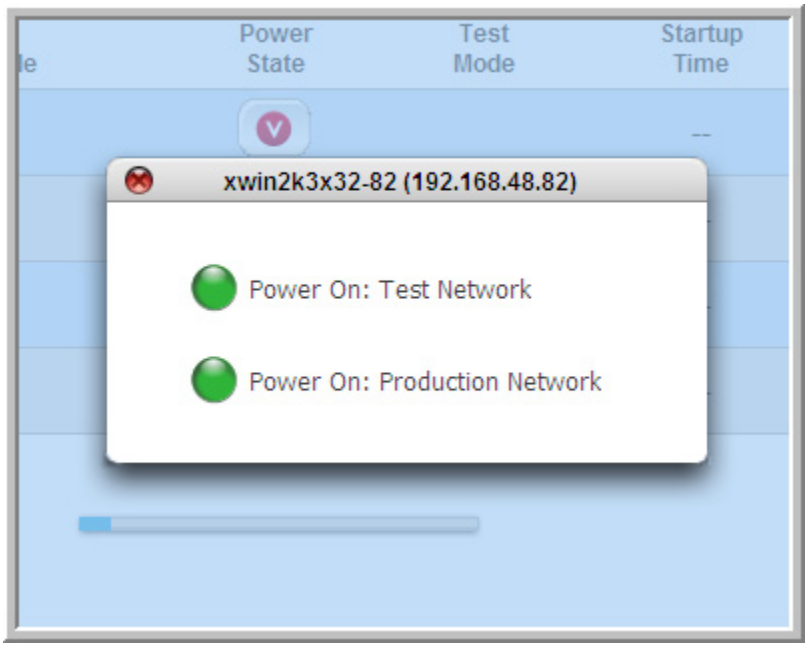
   The startup order will take into account the order you specified in "Configure startup dependencies" on page 264.

To start an *individual* recovery node in *production mode*:

You might be prompted to install application-specific software licenses upon startup of recovery nodes.

1. (Optional) Consider suspending RN tests to “free up” resources.
2. (Redhat 5.x RNs) Install kernel-xen RPM package.
3. From the DR Appliance, go to **DASHBOARD** tab > **RECOVERY NODES** page.
4. Locate the **Power State** column.

   When you click the button in the **Power State** column, a box appears that offers you a choice between the following two alternatives: **Power On: Test Network** or **Power On: Production Network**.
5. Choose the **Power On: Production Network** power option.



Related Topics

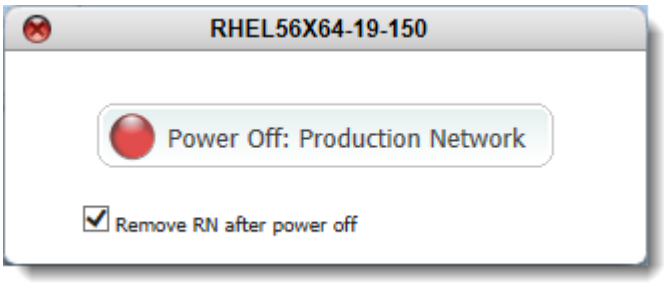
Cancel or override RN startup process

11.9 Stop recovery nodes

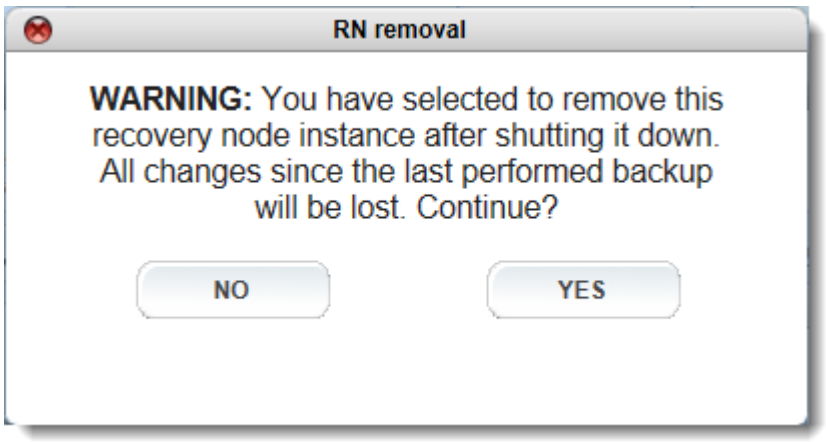
Use this procedure if you are recovering from a failed PN. After you fix the PN, you must stop the recovery node before bringing the PN online; otherwise, there will be network conflicts.

To stop an individual recovery node in *production mode:*

1. Log on to the HA Appliance's onQ Portal. In the event of a site failure, this HA Appliance was formerly your DR Appliance before you changed its role.
2. Go to **DASHBOARD** tab > **RECOVERY NODES** page.
3. Unlock the page.
4. Click the button in the **Power State** column, then **Power Off: Production Network** button, leaving the **Remove RN after power off** check box selected.



5. If you receive the following warning, click **Yes**, then perform an immediate backup.



This warning means that changes to the RN since the last performed back have not been written to the Repository.

Note: If you were testing an RN in production, you'd want to clear the **Remove RN after power off** check box so as to save the RN for future use. However, you're trying to recover from a failed PN, and so it's not important that you keep this RN.

6. Do one of the following:
   - If you are performing this procedure as part of a site disaster, return to (Workflow) Fail back DR to HA Appliance.
   - Otherwise, perform a BMR using QUARK to restore to this RN (aka Recovery PN) as outlined in (Workflow) Fail back an RN to a PN.

11.10 Back up and restore Oracle 11g database on Linux

Oracle recommends that you use RMAN to perform backups and restores of your databases.

Presumably you're already running RMAN scripts to back up and restore your Oracle databases in production—a best practice for any business continuity plan. This procedure walks you how to back up and restore using Quorum's example scripts. Feel free to insert hooks into these backup and restore scripts so as to execute the scripts you're already using.

Quorum's `pre_snap.sh` script is optional; for example, you might prefer to use cron jobs or other scheduler; as such, discuss these options with your Linux Oracle Database Administrator.

You should not have any difficulties using your existing scripts in an onQ environment because backup and restore operations are independent of onQ. However, before you use your scripts in an onQ environment, ensure that they are performing database backups and restores as expected.

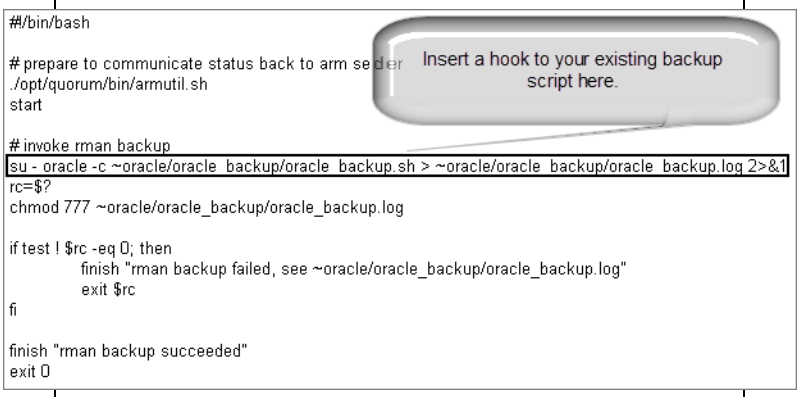
If you do not currently have a backup and restore strategy in place, involve an experienced Linux Oracle Database Administrator before you deploy onQ. onQ is one component of your backup and restore strategy.

(Step 1) To install and modify example RMAN scripts:

This procedure works for Linux 6.2 RNs, though this procedure should also work on other os versions. Quorum's example RMAN scripts use RMAN to back up and restore a single DB instance and uses onQ to activate the backup script. Modify these scripts as needed. The following example is a basic proof-of-concept.

1. Download Quorum's example scripts, then modify these scripts to fit your Oracle environment (RMAN version):

`pre_snap.sh`	Performs the database backup by calling `oracle_backup.sh`. There are many ways to trigger DB backup script: cron jobs, third-party scheduler, or application, for example; however, onQ Portal provides a way to trigger this `pre_snap.sh` script before it backs up the protected volumes (see Step 5). This script can be modified to simply call an existing backup script that you're currently using in your environment:



oracle_backup.sh	Coordinates rman backup of the Oracle database.The onQ Service runs this script before the Quorum backup process. Modify this script to specify the details of your Oracle installation: ORACLE_BASE=**/opt/oracle** ORACLE_HOME=**$ORACLE_BASE/product/11.2.0/db_1** Change the parameters for your database: BACKUP_BASE=**/backup01** BACKUP_AREA=**${BACKUP_BASE}/rdbms/physical/orcl** BACKUP_LOG_DIR=**~/oracle_backup** Modify for your Oracle SID: db=**orcl** This script can be replaced with an existing database backup script that you're currently using in your environment.

`oracle_restore.sh`	Coordinates rman restore/recovery of the Oracle database. Use this script to restore and recover the Oracle database if automated crash recovery fails. Modify this script to specify the details of your Oracle installation: `ORACLE_BASE=`**`/opt/oracle`** `ORACLE_HOME=`**`$ORACLE_BASE/product/11.2.0/db_1`** Change the parameters for your database: `BACKUP_BASE=`**`/backup01`** `BACKUP_AREA=`**`${BACKUP_BASE}/rdbms/physical/orcl`** `BACKUP_LOG_DIR=`**`~/oracle_backup`** Modify for your Oracle SID: `db=`**`orcl`** This script can be replaced with an existing database backup script that you're currently using in your environment. This script should be available in advance of your need to recover: as a best practice, always test a recovery before the need arises.

`armutil.sh`
`armSender`
`vmmonCtl`
	Communicates `pre_snap.sh` execution status to onQ, and depends on `armSender` and `vmmonCtl` executables in `x64` or `i386` folder.

2. Log on to the PN as oracle user (for example, **`su - oracle`**) and create a `/home/oracle/oracle_backup` directory.
3. Now that you've modified `oracle_backup.sh` and `oracle_restore.sh`, copy them to `/home/oracle/oracle_backup` on the PN.
4. As `root` user, copy `armutil.sh` and `armSender` and `vmmonCtl` from x64 or i386 folder, depending on the PN's architecture, to `/opt/quorum/bin` on the Linux PN that's running the Oracle database.

The `armSender` and `vmmonCtl` are tools to communicate status to the onQ Manager.

5. As `root` user, copy `pre_snap.sh` to `/opt/quorum/bin/` on the PN, then activate this script on the onQ Portal.
   a. Log on to the HA Appliance's onQ Portal.
   b. From the **PROTECTION CONFIG** tab, select the PN that's running the Oracle 11g database, then **MODIFY** > **ADVANCED**.
   c. Enable the pre snapshot script on that PN by selecting the **Execute Pre Snapshot script?: Yes** radio button.

That's it! Now you're ready to test the backup process.

(Step 2) To back up the Oracle database:

This procedure works for Linux 6.2 RNs, though this procedure should also work on other os versions.

1. Load some known data into the Oracle database (for example in ORCL catalog).
2. Initiate a backup, or wait for a scheduled backup, of the protected node.

   The PN's Event Log displays the following messages as onQ builds the RN. These messages also appear in `oracle_backup.log` on the PN:
   - Pre-snapshot command returned: rman backup succeeded (as coded in `pre_snap.sh`), if rman backup passes. For example:

```
RHEL62x64-19-87: system: 2014-12-30 13:34:19: Pre-
snapshot command returned: rman backup succeeded
```

   - Pre-snapshot command returned: rman backup failed (as coded in `pre_snap.sh`), if rman backup fails. For example:

```
RHEL62x64-19-87: system: 2014-12-30 13:11:37: Pre-
snapshot command returned: rman backup failed, see
~oracle/oracle_backup/oracle_backup.log
```

   If the RMAN backup fails, consult Oracle's RMAN documentation to fix the problem. The RMAN backup must succeed before you proceed with the restore.

If the backup succeeded, you're ready to test the restore process.

(Step 3) To restore the Oracle database:

This procedure works for Linux 6.2 RNs, though this procedure should also work on other os versions.

1. Start the Recovery Node in test mode, after the RN build completes.

   After the RN boots, the Oracle database should perform crash recovery automatically. The database either opens successfully or it doesn't:

   - (Case A) The database opens successfully and is ready for connection with no need for recovery, as shown in the following output; therefore, skip to Step 4.

```
[root@RHEL62x64-19-87 oracle_backup]# su - oracle
[oracle@RHEL62x64-19-87 ~]$ sqlplus / as sysdba
SQL*Plus: Release 11.2.0.1.0 Production on Tue Dec
23 09:22:06 2014
Copyright (c) 1982, 2009, Oracle.  All rights
reserved.

Connected to an idle instance.
SQL> startup;
ORACLE instance started.

Total System Global Area 1603411968 bytes
Fixed Size                  2213776 bytes
Variable Size            1124075632 bytes
Database Buffers          469762048 bytes
Redo Buffers                7360512 bytes
Database mounted.
Database opened.
SQL> connect SCOTT/<password>
Connected.
SQL> select count(*) from qa_test;
   COUNT(*)
   ----------
   27472
```

   - (Case B) The Oracle database does not open successfully, as shown in the following output. You need to perform the recovery

procedure; therefore, proceed to the next step:

```
[root@RHEL62x64-19-87 oracle_backup]# su - oracle
[oracle@RHEL62x64-19-87 ~]$ sqlplus / as sysdba
SQL*Plus: Release 11.2.0.1.0 Production on Tue Dec
23 09:22:06 2014
Copyright (c) 1982, 2009, Oracle.  All rights
reserved.

Connected to an idle instance.
SQL> startup;
ORACLE instance started.

Total System Global Area 1603411968 bytes
Fixed Size                  2213776 bytes
Variable Size            1124075632 bytes
Database Buffers          469762048 bytes
Redo Buffers                7360512 bytes
Database mounted.
ORA-03113: en-of-life on communication channel
Process ID: 2235
Session ID: 1 Serial number:
```

2. On the RN, run the restore/recovery script to restore the Oracle database onto the RN:

```
# cd /home/oracle/oracle_backup
# ./oracle_restore.sh
```

Optionally, you can run the commands manually as shown in the following example output:

Invoke the restore/recovery script:

```
[oracle@RHEL62x64-19-87 ~]$ cd
/home/oracle/oracle_backup
[oracle@RHEL62x64-19-87 oracle_backup]$
./oracle_restore.sh
Recovery Manager: Release 11.2.0.1.0 - Production on
Tue Dec 23 09:32:15 2014

Copyright (c) 1982, 2009, Oracle and/or its
affiliates.  All rights reserved.

RMAN>          set echo on;
2>          connect target *;
3>          shutdown immediate;
4>          startup mount;
5>          restore database;
6>          shutdown immediate;
7>          startup mount;
8>          recover database;
9>          alter database open;
10>          exit;
echo set on

connected to target database: ORCL (not mounted)

using target database control file instead of
recovery catalog
Oracle instance shut down

connected to target database (not started)
Oracle instance started
database mounted

Total System Global Area    1603411968 bytes
Fixed Size                     2213776 bytes
Variable Size               1124075632 bytes
Database Buffers             469762048 bytes
Redo Buffers                   7360512 bytes
...
```

Restore/recovery begins and completes:

```
Starting restore at 23-DEC-14
allocated channel: ORA_DISK_1
...
restore from backup set
...
channel ORA_DISK_1: restore complete, elapsed time:
00:02:45

Finished restore at 23-DEC-14
database dismounted
Oracle instance shut down

connected to target database (not started)
Oracle instance started
database mounted

Total System Global Area    1603411968 bytes
Fixed Size                     2213776 bytes
Variable Size               1124075632 bytes
Database Buffers             469762048 bytes
Redo Buffers                   7360512 bytes
Starting recover at 23-DEC-14
allocated channel: ORA_DISK_1
channel ORA_DISK_1: SID=18 device type=DISK

starting media recovery
media recovery complete, elapsed time: 00:00:03
Finished recover at 23-DEC-14
database opened
Recovery Manager complete.
```

3. Verify that the RN's Oracle database is up and running and in a usable state, and that it shows the expected data:

```
[oracle@RHEL62x64-19-87 oracle_backup]$ sqlplus /
as sysdba

SQL*Plus: Release 11.2.0.1.0 Production on Tue Dec
23 09:44:55 2014
Copyright (c) 1982, 2009, Oracle.  All rights
reserved.

Connected to:
Oracle Database 11g Enterprise Edition Release
11.2.0.1.0 - 64bit Production
With the Partitioning, OLAP, Data Mining and Real
Application Testing options

SQL> connect SCOTT/<password> ' Able to connect to
DB after restore and recovery
Connected.
SQL> select count(*) from qa_test;
COUNT(*)
--------
   27472

SQL>EXIT;
```

4. Start the Recovery Node in production mode; when prompted, do not remove the RN.

Related Topics

Run custom backup scripts

Back up and restore Oracle 10g+ database on Windows

11.11 Back up and restore Oracle 10g+ database on Windows

The following procedures walks you through how to back up and restore an Oracle 10g+ database on a Windows Recovery Node. These procedures assume that you are knowledgeable and comfortable with Oracle databases.

Restoring and recovering an Oracle database, after you have booted the recovery node (RN), is a unique process from that of other databases. These procedures cover the prerequisites for hot backups with Volume Shadow Copy, and the required tools developed by Quorum to recover and open your Oracle Database on your recovery node.

These procedures outline the steps to recover an Oracle database with `ARCHIVELOG` mode enabled and the Oracle VSS Writer service installed. Hot backups are not possible with an Oracle database using `NOARCHIVELOG` mode. Databases using `NOARCHIVELOG` mode must be put in a consistent (shutdown) state in order to perform a backup.

Recovery, maintenance and troubleshooting of Oracle databases on the protected node (PN) will not be covered in these procedures.

These procedures might not work with your specific configuration; as such, discuss your configuration with Quorum Support before you perform this procedure.

(Step 1) To prepare your environment:

1. Verify that your RN is running:
   - Windows 2003 or greater
   - Oracle Database 10g or greater
2. Ensure that you have your archive logs stored on a separate volume from your database (DBF) files.

   Oracle requires that you put archive logs on a separate volume when performing backups with the Oracle VSS Writer service.

   Quorum cannot recommend an architecture that violates Oracle's specifications. However, it should be noted that Quorum did not observe any adverse effects when putting the database and archive logs on the same volume while performing VSS backups and restores.

3. From the onQ Portal's **Protection Config** tab, verify that you have the following volumes protected.
   - The volume that contains your database files.
   - The volume that contains the archive logs.
4. For the protected node on which you want to perform hot backups and for every Oracle database instance, install the Oracle VSS Writer service (`service.msc`).

   The Oracle VSS Writer service is installed by default with Oracle Database 11g. Oracle Database 10g releases require a manual installation of the Oracle VSS Writer service. The installation instructions for the Oracle VSS Writer service can be found here.
5. Configure the Oracle services:
   a. Make sure Oracle is in `ARCHIVELOG` mode.
   b. Configure **Oracle VSS Writer** service with the `SYSDBA` user, mode `Automatic`; then, test that the user account has `SYSDBA` set correctly:

```
> sqlplus user_account as SYSDBA
```

   For example:

   Where `sysadm_amy` is the *user_account*.

```
> sqlplus sysadm_amy as SYSDBA
```

   c. Configure the **Volume Shadow Copy Service** with `Local Service` account, mode `Automatic`.
   d. Configure the **Software Shadow Copy Provider** service with `Local Service` account, mode `Automatic`.

(Step 2) To back up the Oracle database:

Perform the following procedure on the PN and for each database that you want to restore.

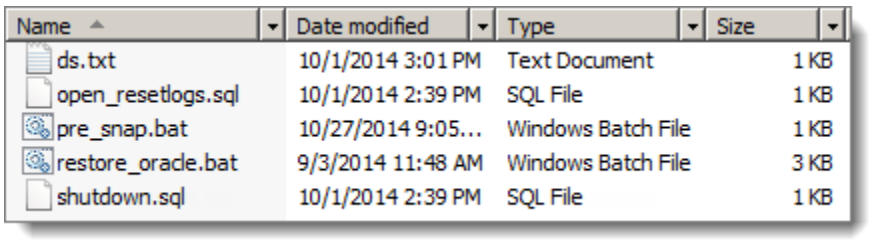
1. Verify that the Oracle VSS Writer service is `running`.

```
> oravssw /q /status
OracleVssWriterORCL - The service is running.
```

2. On the PN, verify that the VSS writers are `Stable`.

```
> vssadmin list writers
…
Writer name: 'Task Scheduler Writer'…
   State: [1] Stable
   Last error: No error
Writer name: 'VSS Metadata Store Writer'…
   State: [1] Stable
   Last error: No error
Writer name: 'Performance Counters Writer'…
   State: [1] Stable
   Last error: No error
Writer name: 'Oracle VSS Writer - ORCL'…
   State: [1] Stable
   Last error: No error
Writer name: 'System Writer'…
   State: [1] Stable
   Last error: No error
Writer name: 'ASR Writer'…
   State: [1] Stable
   Last error: No error
Writer name: 'Registry Writer'…
   State: [1] Stable
   Last error: No error
Writer name: 'COM+ REGDB Writer'…
   State: [1] Stable
   Last error: No error
Writer name: 'Shadow Copy Optimization Writer'…
   State: [1] Stable
   Last error: No error
Writer name: 'WMI Writer'…
   State: [1] Stable
   Last error: No error
```

3. On the PN, unzip the Quorum Oracle database recovery scripts in the `C:\Program Files\Quorum\usr` folder.
4. Modify the `restore_oracle.bat` batch file:



Name	Date modified	Type	Size
ds.txt	10/1/2014 3:01 PM	Text Document	1 KB
open_resetlogs.sql	10/1/2014 2:39 PM	SQL File	1 KB
pre_snap.bat	10/27/2014 9:05...	Windows Batch File	1 KB
restore_oracle.bat	9/3/2014 11:48 AM	Windows Batch File	3 KB
shutdown.sql	10/1/2014 2:39 PM	SQL File	1 KB

   a. Make a backup copy of the `restore_oracle.bat` batch file.
   b. Change the placeholder database names to match your Oracle database names. Each instance should be separated by a single space and encapsulated by parenthesis.

```
> set DB_INSTANCES=(database_instance_name)
```

   For example:

   Where there are three database instances: `orcl`, `orcl1`, and `orcl2`:

```
> set DB_INSTANCES=(orcl oracl1 oracl2)
```

   c. Save your changes.
5. Initiate a backup of the PN, or wait for a scheduled backup to complete.
6. Repeat Step 2: Verify that the VSS writers are `Stable`.
7. Verify that the backup completed successfully. Go to Monitor backups.
8. Check the PN's Event Log for VSS-related errors. If there are errors, resolve them. For troubleshooting information, see the appropriate return code in A0801.

(Step 3) To restore the Oracle database:

After the RN boots, you can now begin the Oracle Database recovery process. This process can be performed with an RN booted in test mode or production mode. Perform the following procedure on the RN and for each database instance that you want to restore. This procedure includes two ways to restore your database(s). Begin with the recommended solution.

(Recommended) Using restore_oracle.bat

This procedure uses Quorum's `restore_oracle.bat` scripts. This solution is recommended.

1. With the RN, in test or production mode, navigate to the `C:\Program Files\Quorum\usr` folder, then run the `restore_oracle.bat` file completely.

   You will be prompted for your database password twice for each database instance: once for a shutdown command and once for an open **`resetlogs`** command.

2. Verify that your database(s) opens without any errors.
   - If the database(s) opens without error, you've successfully restored your database(s). You're done!
   - If the database(s) fails to open successfully, the cause is the result of known bugs in Oracle's VSS Writer. Perform the troubleshooting procedure outlined in (Step 4, if necessary) To fix an incomplete database recovery:.

Using "manual" process

This procedure uses a manual process to restore your database(s). This solution is not automated as is the case with using Quorum's `restore_oracle.bat` scripts, which executes the same commands in the following manual process.

1. Shut down the database that you want to restore.

```
> set ORACLE_SID=database_instance_name
> sqlplus SYS as SYSDBA
SQL> SHUTDOWN IMMEDIATE;
SQL> EXIT;
```

   For example:

   Where `orcl` is the *database_instance_name* or Oracle System ID (SID).

```
> set ORACLE_SID=orcl
> sqlplus SYS as SYSDBA
SQL> SHUTDOWN IMMEDIATE;
SQL> EXIT;
```

2. Restore the Oracle VSS Writer's metadata and generate the required archive log, if any.

```
> qvss.exe restore
```

3. Open the database for service and reset the redo logs.

```
> sqlplus SYS as SYSDBA
SQL> ALTER DATABASE open resetlogs;
SQL> EXIT;
```

4. Compare the PN's and RN's row counts. They should match, if no changes were made after the backup.

```
> sqlplus system@database_instance_name
SQL> select count (*) from table_name;
```

For example:

Where `testtable` is the *table_name* and where `orcl` is the *database_instance_name* or Oracle System ID (SID).

```
> sqlplus system@orcl
SQL> select count (*) from testtable;
```

5. Check the PN's Event Log for VSS-related errors. If there are errors, resolve them. For troubleshooting information, see the appropriate return code in A0801.

6. Verify that your database(s) opens without any errors.
   - If the database(s) opens without error, you've successfully restored your database(s). You're done!
   - If the database(s) fails to open successfully, the cause is the result of known bugs in Oracle's VSS Writer. Perform the troubleshooting procedure outlined in (Step 4, if necessary) To fix an incomplete database recovery:.

(Step 4, if necessary) To fix an incomplete database recovery:

If you attempted to restore an Oracle database either manually or using `restore_oracle.bat` and your database(s) fails to open, you have an incomplete database recovery. This procedure includes two ways to fix an incomplete database recovery. Begin with the recommended solution.

(Recommended) Using "recover database until cancel" method

This solution uses Oracle's *recover database until cancel* recovery method to fix an incomplete database recovery. This procedure is the easiest way to fix such a problem, and is recommended before attempting any other incomplete database recovery method.

1. Restart a *fresh* RN, in test or production mode. If you're using `restore_oracle.bat`, execute that file now, but choose not to start the databases at the end of batch file.
2. Set the `ORACLE_SID` to the database SID. Where the SID is `ORCL`:

```
C:\> SET ORACLE_SID=ORCL
```

3. Start the sqlplus shell:

```
C:\> SQLPLUS / AS SYSDBA
```

4. Perform the following Oracle DB incomplete recovery procedure, which uses the *recover database until cancel* method:
   a. Execute the *recover database until cancel* method.
   b. Verify that the suggested log file is available or specify an available archived log file to be restored, then press **Enter**.

If the suggested log file is unavailable or there are no archive log files, type **CANCEL**.

```
SQL> RECOVER DATABASE UNTIL CANCEL USING BACKUP
CONTROLFILE;
ORA-00279: change 5619127 generated at 11/18/2014
13:03:07 needed for thread 1
ORA-00289: suggestion :
E:\LOGS\ARCH_863902564_1_81.ARC
ORA-00280: change 5619127 for thread 1 is in se-
quence #81


Specify log: {<RET>=suggested | filename | AUTO |
CANCEL}


++++++++++++++++++++++++++++++++++++++++++++++++++
++++++++++
Specify log: {<RET>=suggested | filename | AUTO |
CANCEL}
CANCEL
Media recovery cancelled.
SQL>
```

5. Open the database with resetlogs:

```
SQL> alter database open resetlogs;
```

6. Verify that the database opens successfully. The database has a new change number in the database view.

```
SQL> SELECT CURRENT_SCN FROM V$DATABASE;
```

If the database fails to open successfully, the cause is a known bug (race condition) in Oracle's VSS Writer. Perform the procedure outlined in Using point-in-time recovery method.

Using point-in-time recovery method

This solution uses Oracle's *point-in-time* recovery method to restore the database to the target time (or SCN) for recovery. This procedure is guaranteed to succeed, but is more complicated; for this reason, use it as a

last resort. Perform the following procedure for each database instance.

1. Restart a *fresh* RN, in test or production mode. If you're using `restore_oracle.bat`, execute that file now, but choose not to start the databases at the end of batch file.
2. Set the `ORACLE_SID` to the database SID. Where the SID is `ORCL`:

```
C:\> SET ORACLE_SID=ORCL
```

3. Start the sqlplus shell:

```
C:\> SQLPLUS / AS SYSDBA
```

4. Perform the following Oracle DB incomplete recovery procedure, which uses the *point-in-time* recovery method.
   a. Determine the missing last change number and record the last `NEXT_CHANGE#` value:

```
SQL> SELECT MAX(NEXT_CHANGE#) FROM V$LOG_HISTORY;
```

   Let's assume the last change value is `118009`.

   b. Recover to the last change number in the last archived redo log on the RN. This value is the last `NEXT_CHANGE#` value (see Step 4) minus 1. Where the last change value is `118009`, therefore, the value is `118008` (118009 -1).

```
SQL> RECOVER DATABASE UNTIL CHANGE 118008 USING
BACKUP CONTROLFILE;
```

5. Open the database with resetlogs:

```
SQL> alter database open resetlogs;
```

6. Verify that the database opens successfully. The database has a new change number in database view.

```
SQL> SELECT CURRENT_SCN FROM V$DATABASE;
```

   If your database still fails to open successfully, contact Quorum Support.

Related Topics

Run custom backup scripts
Back up and restore Oracle 11g database on Linux

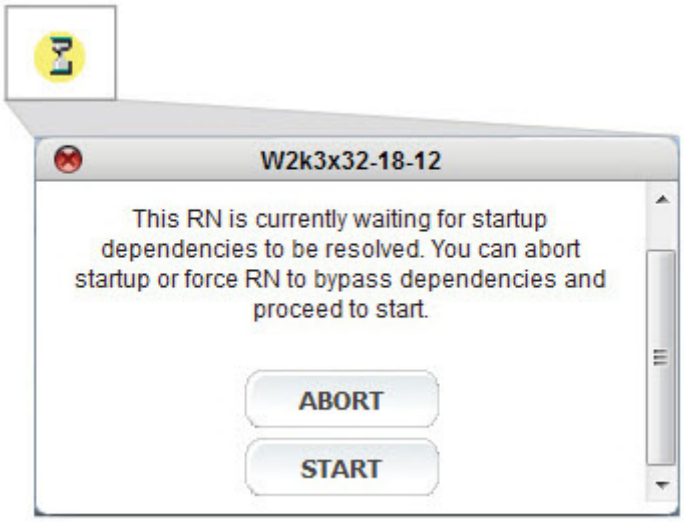
11.12 Cancel or override RN startup process

During the startup process for an RN configured with dependencies, onQ might be waiting for a dependency to become available or for the delay period.

If you are waiting for a dependent RN/server that has no chance of coming online and/or you never specified a timeout, you might need to force or cancel the startup process.

To cancel or override the process:

1. Log on to the onQ Appliance's onQ Portal.
2. Go to **DASHBOARD** tab > **RECOVERY NODES** page.
3. Unlock the page.
4. Click the button in the **Power State** column, then click:
   - **ABORT** - to cancel the startup process. Fix the dependency problem, then restart.
   - **START** - to override the dependencies and start the RN. Fix the dependency problem.



11.13 Launch recovery node's console

Use this procedure if you'd like to test a recovery node.

To launch a recovery node's console:

1. If the RN is an HA recovery node, start it in test mode.

Note: You can use the console in either test mode or production mode.

2. Go to **DASHBOARD** tab > **RECOVERY NODES** page.
3. Click on the **Console** icon. The console launches as a HTML5 web-based VNC console or a Java-based console, depending on your browser version.

If your browser supports websockets and the canvas element, a VNC console opens in a new browser window or tab; however, if you invoke the console in a browser that does not offer this support, the browser reverts to the Java-based console.

- **Web-based VNC Console**. Add an exception to the pop-up blocker as most browsers will block this new window by default.
- **Java Console**. Choose **Open with Java Web Start Launcher** to open the.`jnlp` file. If prompted, click **OK** to accept the security warning. If your desktop cannot recognize this file, then you don't have JRE (Java Runtime Environment) installed. Download the software from the Java site.

Note: If the onQ portal returned the following message, the port required for the VNC session is not available from the hypervisor: `The VNC session is not currently available. Please try again`. During RN power, the port hasn't been assigned yet. During RN power down, the port has already been closed. Simply wait a few seconds, then try again.

4. From the **onQ RN Console**, click **Send Ctrl-Alt-Delete** in the main menu.

   Now you can log on using the host's credentials.

Related Topics

Test Mode
Start recovery nodes on HA Appliance

11.14 Test Mode

When you start a recovery node, you have an opportunity to run the recovery node in test mode. You can only access a recovery node in test mode through the integrated console.

Running a recovery node in test mode means that it runs in a limited private network environment isolated from your production facilities.

The reason for running in test mode is to ensure, without risk to ongoing operations, that the recovery node is intact and will run in the event of a disaster. Keep in mind that this capability is in addition to automatic testing.

The onQ Appliance itself occasionally and briefly runs a recovery node in test mode as part of an internal quality-assurance process. However, best practice is to check your recovery nodes periodically, just as you would ordinary backups.

While an RN is running in test mode the corresponding PN also continues to run although the backup process is suspended. Backups (and subsequent rebuilds of the RN) resume according to current settings when the test mode has finished.

Related Topics

Configure automatic testing of RNs
Start node protection

11.15 Configure automatic testing of RNs

onQ can automatically test (also called *Self Test*) recovery nodes on a predefined schedule. These tests determine if the nodes boot properly and can communicate with the onQ Service. Self Test works on both the HA Appliance and the DR Appliance and even when PNs are on different subnets. Upon completion, onQ sends the test results to all subscribers.

onQ can run tests on *individual* nodes or a *group* of nodes. Test a group of nodes if an application depends on one or more services; otherwise, test nodes individually.

To configure an RN test for an *individual* recovery node:

1. From the HA Appliance or the DR Appliance, go to the **SELF TEST** tab > **ADD INDIVIDUAL** button.
2. In the RN Name drop-down list, choose the recovery node from the list.
3. In the Test Schedule drop-down list, choose when you want the test to run, then **SAVE**.

   After Update – Choose this option if you want the test to run after each incremental update. This option is ideal when your backups are scheduled once every 24 hours.

   Hour – If you have multiple incremental backups in a day, choose this option to avoid getting multiple confirmations after each test. More than one test in a day does not provide any additional benefits.

Related Topics

Test Mode
(Agent-based PNs) Methods of Performing Incremental Backups
Initiate immediate backups

11.16 Suspend and resume RN tests

RN tests (aka *Self-Test*), when running, can reduce the memory pool and can interfere with RN testing. The onQ Portal enables you to suspend all RN

tests that are currently running.

Suspending and resuming RN tests do not generate alert emails. Moreover, an alert email will not be generated on the next successful test after RN tests resume. However, an alert will be generated if an RN test failure occurs after RN tests resume.

To suspend all RN tests:

1. Do one of the following:
   - From the HA Appliance or the DR Appliance, go to the **SELF TEST** tab > **ADD INDIVIDUAL** button.
   - Go to **DASHBOARD** tab > **RECOVERY NODES** page, then unlock the page.
2. Click on the **Suspend Self Test** button, then **Suspend**.

   The Suspend icon appears in the page header to the right of the Protection On/Off display, and messages appear in the Self Test log:



`Suspend RN Self-Tests of ALL`	onQ instructs the scheduler to suspend all RN tests.
`Request to stop Self-Test for <host>`	Appears if an RN test is currently running.
`Self-Test stopped`	Appears when the scheduler successfully stops the RN test(s).

To restart all RN tests:

- From the HA Appliance or the DR Appliance, go to the **SELF TEST** tab > **ADD INDIVIDUAL** button.
- Go to **DASHBOARD** tab > **RECOVERY NODES** page, then unlock the page.

3. Click on the **Resume Self Test** button, then **Resume**.

   The Self Test log indicates `Resume RN Self-Tests for ALL`.

11.17 Restore PNs Using QUARK

The following disaster recovery workflows prompt you to use Quorum Ultimate Automated Recovery Kit (QUARK):

- (Workflow) Fail over HA Appliance to DR Appliance
- (Workflow) Fail back DR to HA Appliance
- (Workflow) Fail over a PN to an RN
- (Workflow) Fail back an RN to a PN

QUARK is an easy-to-use wizard that walks you through two types of restores:

- **Bare Metal Restore (BMR)**. You have a server failure or site failure and you want to restore a Recovery PN from scratch.
- **Incremental/Reversion Restore**. You have a server failure and you want to restore to the latest snapshot, or you have a corruption problem and need to restore to a previous snapshot (aka *reversion*).

You can also restore PN data using two other methods: file level restore (FLR) and Windows Share Restore (WSR). These methods are commonly used in non-disaster recovery scenarios.

To learn more about BMR/QUARK support and limitations, go to:

- "Bare Metal Restore Limitations" in onQ Release Notes
- "Bare Metal Restore Support" in onQ Release Notes

(Step 1) To download the QUARK software:

QUARK supports two boot options:

- **USB**. This media is used for most installs. If for any reason QUARK cannot inject the necessary drivers, you can add drivers to this USB for use during the BMR without needing additional media.
- **ISO**. Burn this ISO to a CD-ROM or, in the case of virtualization, save this ISO to your network and point to it from your virtualization infrastructure.

  1. Log on to the onQ Portal.
  2. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **SOFTWARE UPGRADES** page > **Check for Updates** button.

This step ensures that you have the most current version of QUARK for your onQ Appliance.

3. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **DOWNLOADS**, then do one of the following.
   - **USB method**. Select the `Quark_usb_BCV-<release>-<date>-<time>.zip`, **Download**, then save the `.zip` file to a USB key.
   - **ISO method**. Select the `Quark_BCV-<release>-<date>-<time>.iso`, **Download**, then save the `.iso` file to your virtualization infrastructure, or burn it to a CD-ROM.

   Now you're ready to perform your BMR.

(Step 2) To restore data to your PN:

In this procedure, *Recovery PN* (also called *Recovery OS, BMR target*, *restore target*, or simply *target*) refers to the server to which you are restoring. This Recovery PN can be a physical or virtual machine.

If your Recovery PN is a virtual machine, QUARK requires a "shell" onto which QUARK can load its snapshot(s). You can retrieve this "shell" by building a new virtual machine from scratch, restoring to the original PN, using a snapshot of your virtual machine, or deploying a clone from a template. In any case, this virtual machine need not look like your Recovery PN. Simply reconfigure the machine's BIOS to boot QUARK, then QUARK will walk you through the BMR.

Note: Although QUARK attempts to inject the necessary drivers for your system, it's a best practice to have a copy of these drivers as part of your business continuity plan.

Your Recovery PN must have enough resources to accommodate the new data. To be safe, ensure that your Recovery PN has at least the same amount of disk space and memory as the failed PN.

In the event of a restore failure, QUARK provides you the opportunity to resume the restore (see Step 8 and Step 12). However, this capability will not be available if you:

- **Restart protection**. Therefore, do not stop, then start protection (aka restart) on the onQ.

- **Restart a BMR for a different PN from the same onQ**. Therefore, wait for the BMR to complete before beginning another from the same onQ.

1. Verify that your platform is supported. Go to "Support" in onQ Release Notes.
2. Configure Hardware RAID on the Recovery PN.

   If necessary, configure any hardware RAID controllers in the BIOS of the Recovery PN. For instructions, refer to the hardware manufacturer's documentation.
3. Launch the QUARK wizard:

   a. Configure the Recovery PN's BIOS to boot QUARK.

   b. Boot the Recovery PN to load the QUARK image.

   The **Start BMR** page appears.
4. From the **Start BMR** page, allow the QUARK wizard to check for basic network and disk controller drivers on which QUARK depends, then **Next**.

   If the QUARK wizard detects a missing driver, upload the missing network driver.

Note: To manually launch the **Add Drivers** utility, run the following command from the command prompt:

```
# QUARK_Wizard.exe AddDriver
```

5. (Optional) In the **Network Configuration** page, change the default network configuration, then **Next**:
   - If the wizard detects a DHCP server, the wizard displays the assigned default IP address and netmask.
   - If DHCP is not enabled on the network, the wizard displays the assigned default Windows workgroup address.

   a. Select a NIC from the **Interface** drop-down list.

   b. Specify a *temporary* **IP address** for the restore target. DO NOT set the IP address to the PN's original IP address; doing so prevents onQ from communicating with the restore target. After onQ restores PN data, the restore target will automatically be set to the

PN's original IP address. Select the **Set Static IP** check box to set the default IP address to a static IP.

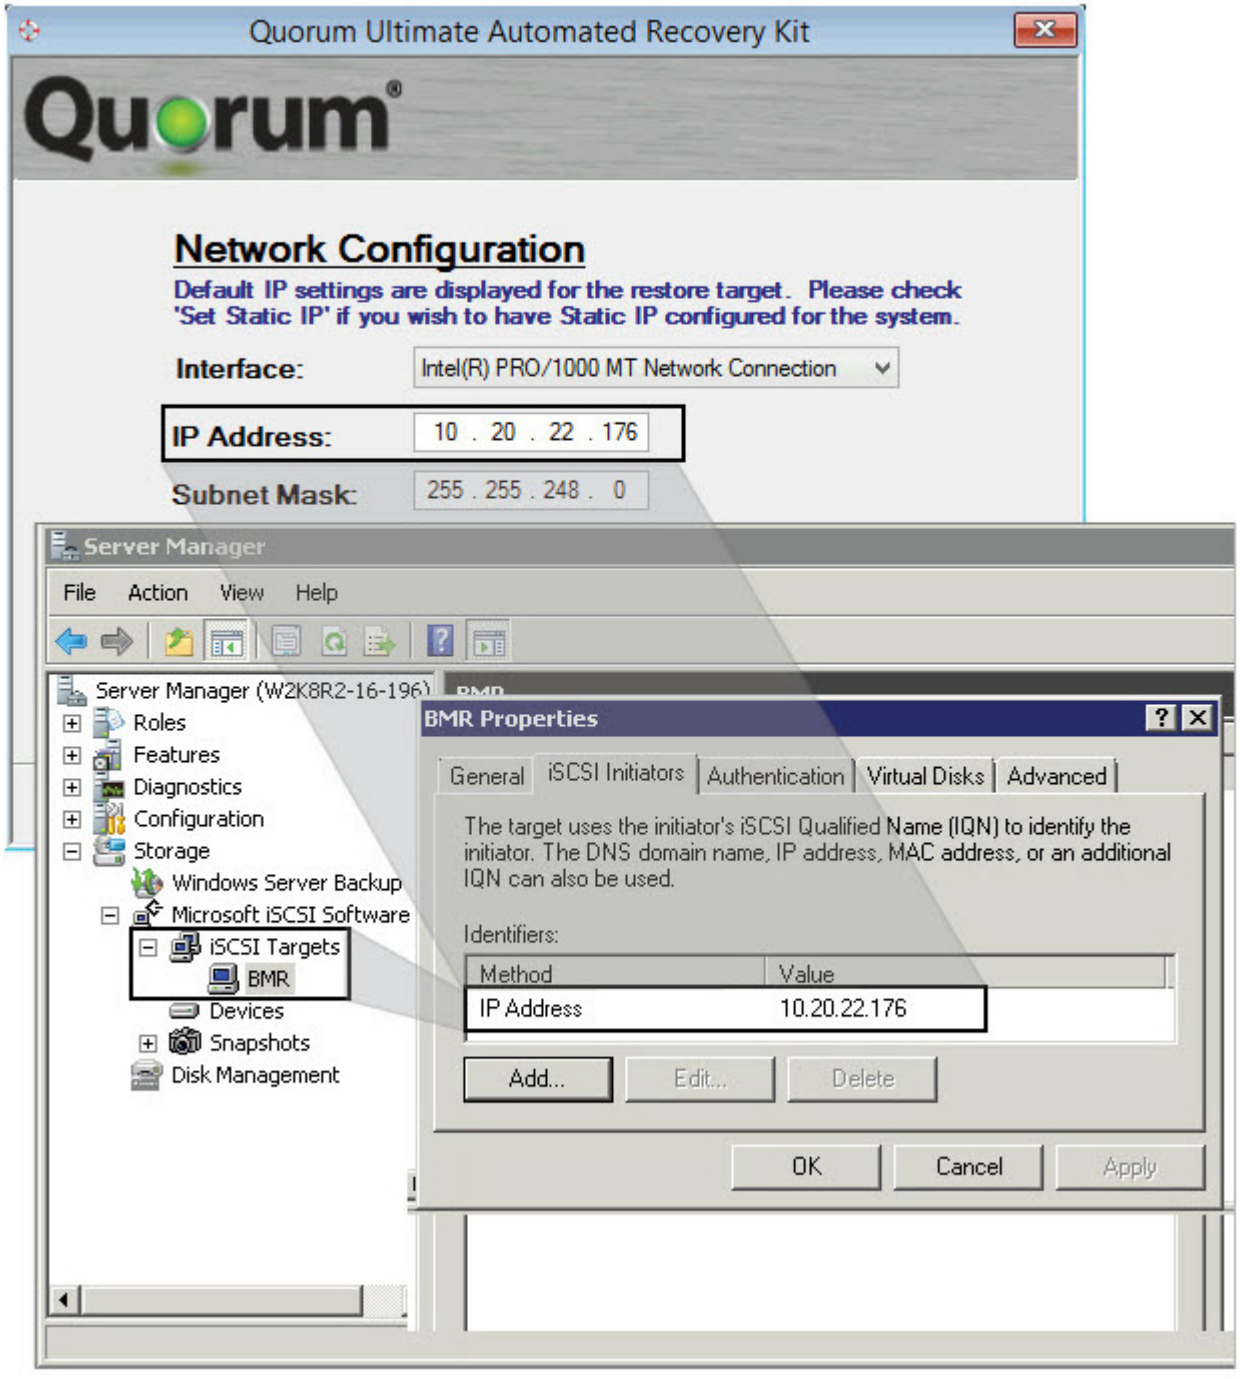
The **Disk Cleanup / iSCSI Configure** page appears.

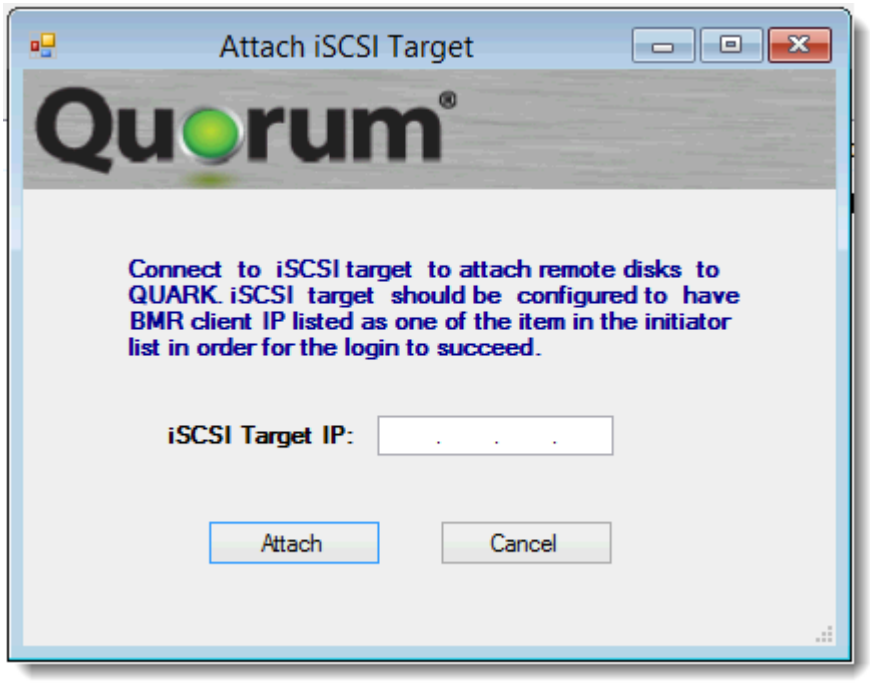
6. Do one of the following:
   - **Reset all local disks**. Select this option if you intend to perform a BMR.

Note: Resetting the disks ensures that you do not write PN data on top of corrupted data. Reformatted and clean disks are ideal. However, if you've already created the disk partitions outside of QUARK using a disk partition tool, you can skip this task.

   - **Skip disk reset**. Select this option if you intend to perform an Incremental/Reversion or your last restore interrupted prematurely.
   - **Configure iSCSI disks**. Select this option if your Recovery PN has iSCSI disks. The **QUARK Disk Configuration** wizard launches. (1) Add QUARK (BMR target) to the iSCSI target's initiator list, then (2) In **Disk Configuration** page, specify the

iSCSI target's IP address to attach the iSCSI devices to QUARK.





The **Log on to onQ** page appears.

7. In the **Log on to onQ** page, type the IP Address and credentials of the onQ Appliance that manages the Recovery PN's snapshot to which you want to restore, then **Next**.

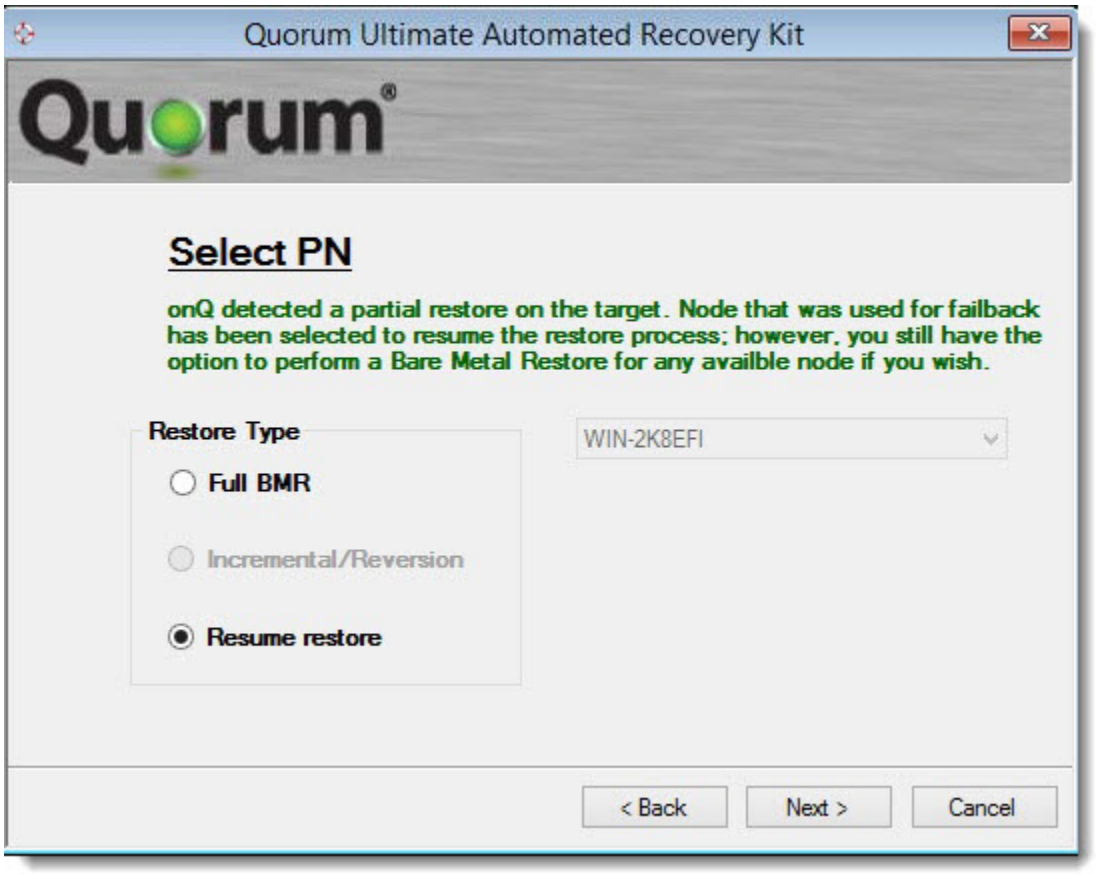
   The onQ Appliance analyzes the snapshots that belong to the Recovery PN, and determines the restore types (**Full BMR** or **Incremental**) that are available for this recovery, then displays the **Select PN** page.

   If this process fails, browse the list of error messages in (Step 4) To interpret QUARK Errors: to correct the problem.

8. In the **Select PN** page, select the restore type and the Recovery PN from the drop-down list of PNs that the onQ Appliance manages, then **Next**.
   - **Full BMR**. Restore a Recovery PN from scratch. If the onQ Appliance doesn't have the incremental backups that are needed to perform the Incremental/Reversion restore (that is, the PN is a new server and does not have any data on its disk), the wizard displays the Full BMR option only; otherwise, both options are available.
   - **Incremental/Reversion**. Use the **Incremental/Reversion** restore

type to: (1) Restore to the latest version; simply choose the latest snapshot, or (2) Revert to a previous version (also known as a *reversion*); simply select the snapshot that represents a point in time. This scenario is ideal when you know that the latest version is corrupt.

- **Resume restore**. Appears when onQ detects an incomplete restore (aka restore failure) on the target. An incomplete restore (see Step 12) may occur if a restore doesn't complete within 48 hours. Choose Resume restore to resume "beginning where QUARK left off." You also have the choice to start over.

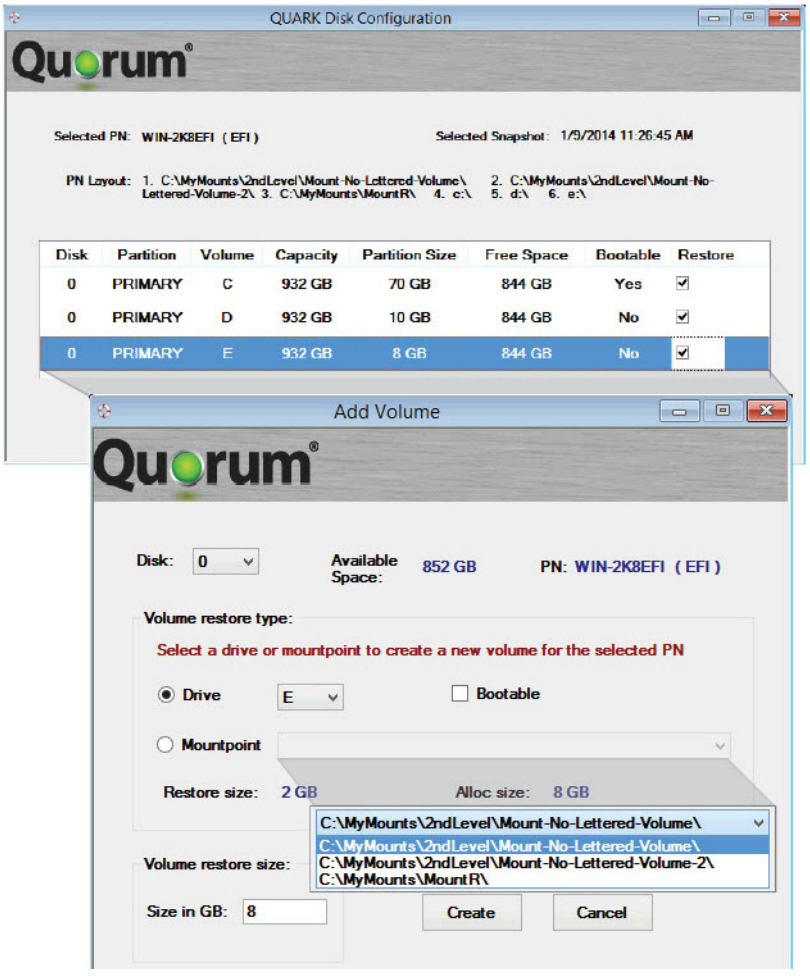


9. In the **Select Snapshot** page, select the snapshot to which you want to restore the Recovery PN, then **Next**.
10. From the **Begin PN Restore** page, click **Next** to confirm your PN and snapshot selection. The **QUARK Disk Configuration** wizard launches.

11. Configure the Recovery PN with the same disk configuration as the original PN, then **Close**.

Note: If a USB/CD-ROM takes a PN drive letter, QUARK automatically reassigns the USB/CD-ROM a different drive letter.

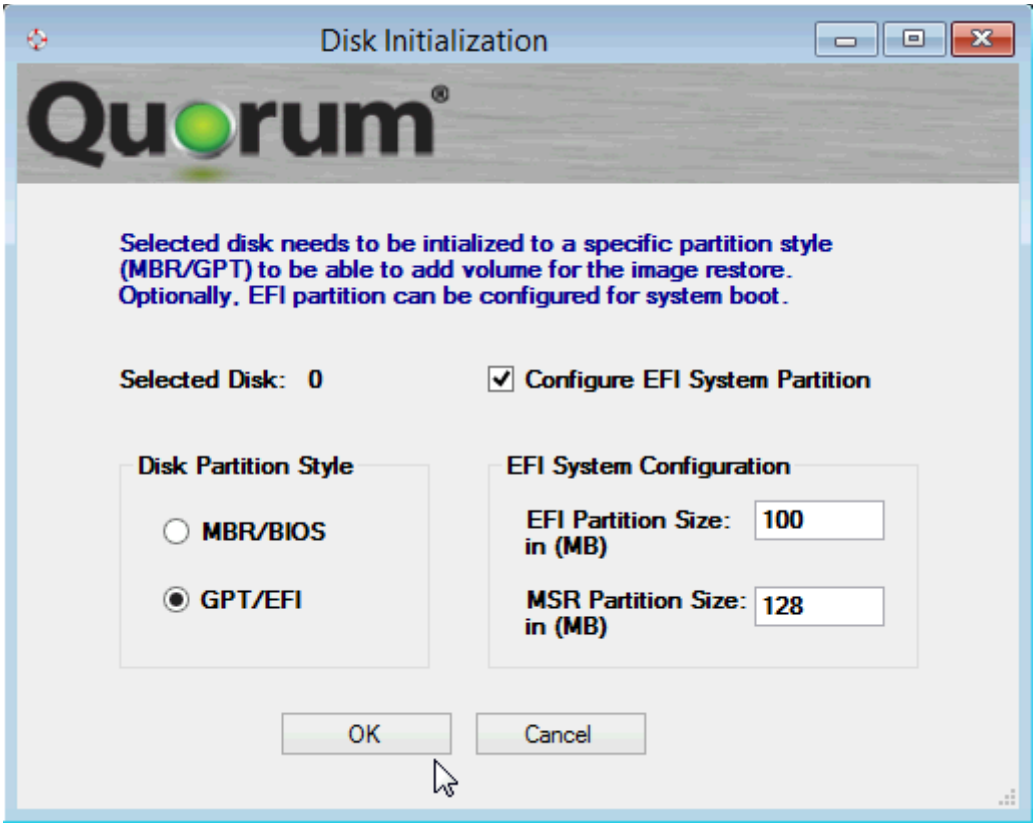
a. For every volume that appears in the PN Layout field, add a corresponding volume. The volume can be either a drive or a mountpoint. Optionally, you can change the **Volume restore size** of the drive or mountpoint; the default is the allocated size as shown in the onQ Portal.



b. If the drive is a boot drive, specify the partition style when the **QUARK Disk Configuration** wizard prompts you:
   - If you want the PN to boot with BIOS, select **MBR/BIOS**.

- If want the PN to boot with UEFI, select **GPT/EFI**.

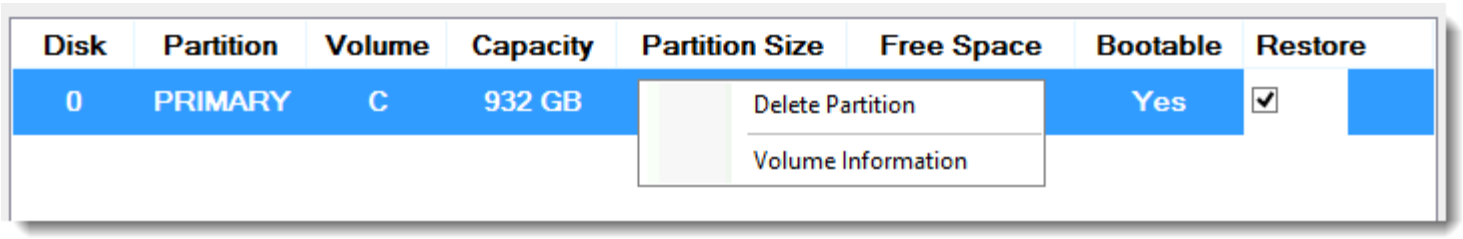
Optionally, change the partition size.



c. Select the volumes that you want to restore. Typically you'll want to select all the volumes; however, you might want to restore only the operating system so as to get the PN up and running quickly as demonstrated in the example below; you can always restore the remaining volumes later using QUARK, FLR, WSR, or a third party tool.

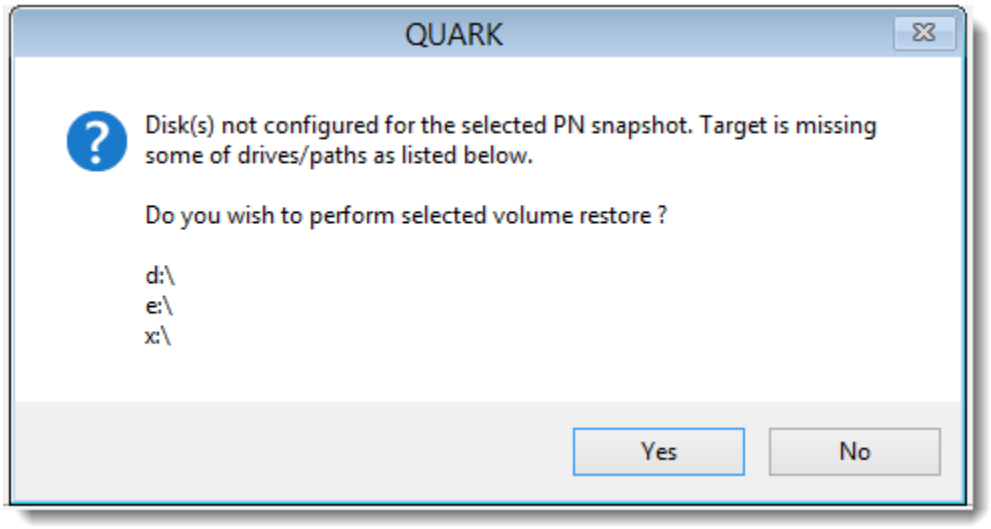
Disk	Partition	Volume	Capacity	Partition Size	Free Space	Bootable	Restore
0	PRIMARY	C	932 GB	70 GB	844 GB	Yes	☑
0	PRIMARY	D	932 GB	10 GB	844 GB	No	☐
0	PRIMARY	E	932 GB	8 GB	844 GB	No	☐

d. If you want to delete a volume or view volume information (for example, file system), right-click on the volume and **Delete Partition** and **Volume Information** respectively.



e. When the **Begin PN Restore** page appears, click **Next**.

If you receive the following error message, the volumes that you added do not match the PN layout. However, as outlined in Step c, you do not need to restore all the volumes.

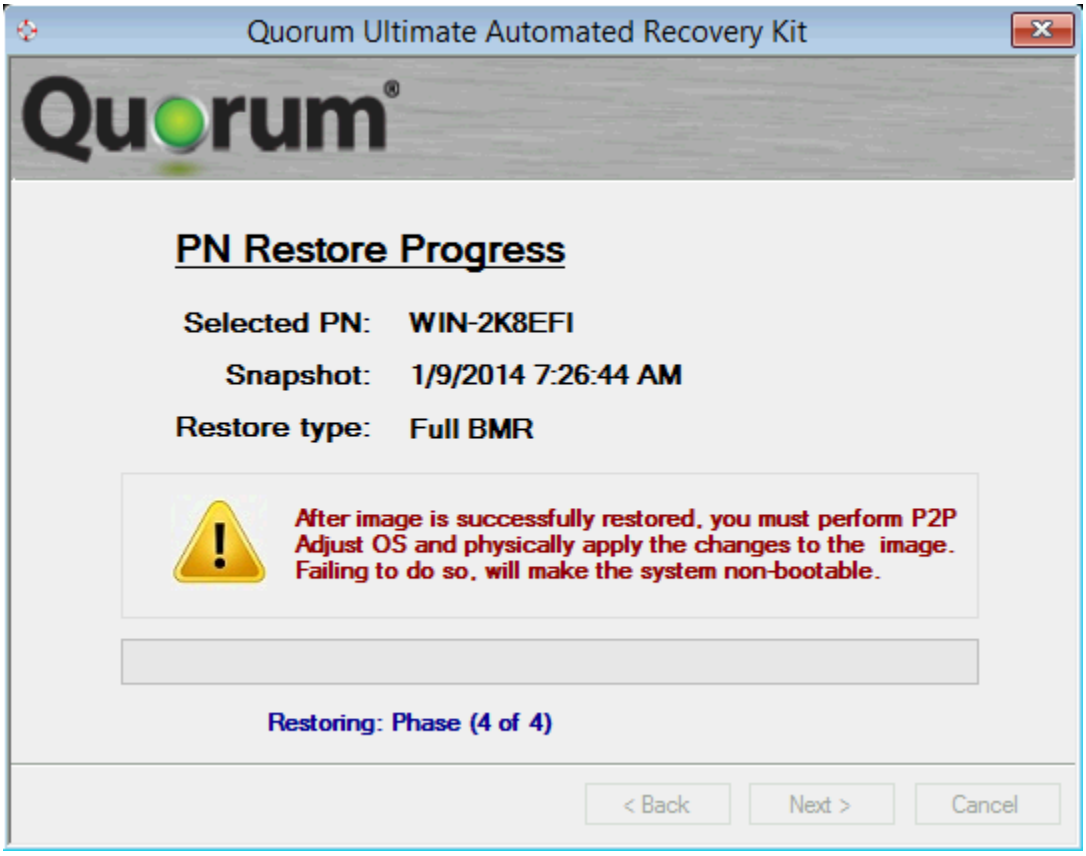


f. Wait while the wizard formats the drives. This process can take several minutes.

Note: If you accidently close the QUARK Disk Configuration wizard or you need to reconfigure your disks before you restore, you can launch the wizard from the command prompt:

```
# QUARK_Wizard.exe showDiskConfig
```

The **PN Restore Progress** page appears.

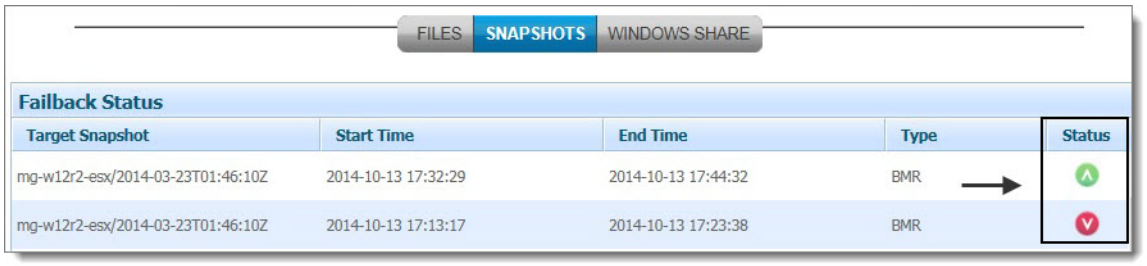


12. Wait from the onQ Appliance to restore the Recovery PN to the snapshot you selected.

a. If the onQ Appliance detects a problem, the QUARK wizard aborts the BMR, then displays a status message. Browse the list of error messages in (Step 4) To interpret QUARK Errors: to correct the problem.

b. If QUARK does not display any errors, verify that the restore completed successfully. Log on to the onQ Appliance's onQ Portal, then go to **RESTORE** > **SNAPSHOTS**. Locate the target snapshot

that you chose. If the **Status** icon is red, hover-over the icon to display the status message.



If the status message indicates `Incomplete: Unexpected Disconnect`, then you will need to launch QUARK and perform the restore again. When you select the PN from **Select PN** page, onQ notifies QUARK of the restore failure, then QUARK provides you the opportunity to resume the restore (see Step 8).

c. If the status message indicates `Completed Successfully` and if you attached iSCSI devices in Step 6, verify that the Microsoft iSCSI initiator assigned the correct drive letters for each iSCSI device. If not, change the drive letters: **Computer Management** > **Disk Management**.

Sometimes the restored PN can have incorrect drive letters for iSCSI devices if the target has an unexpected Microsoft iSCSI initiator version.

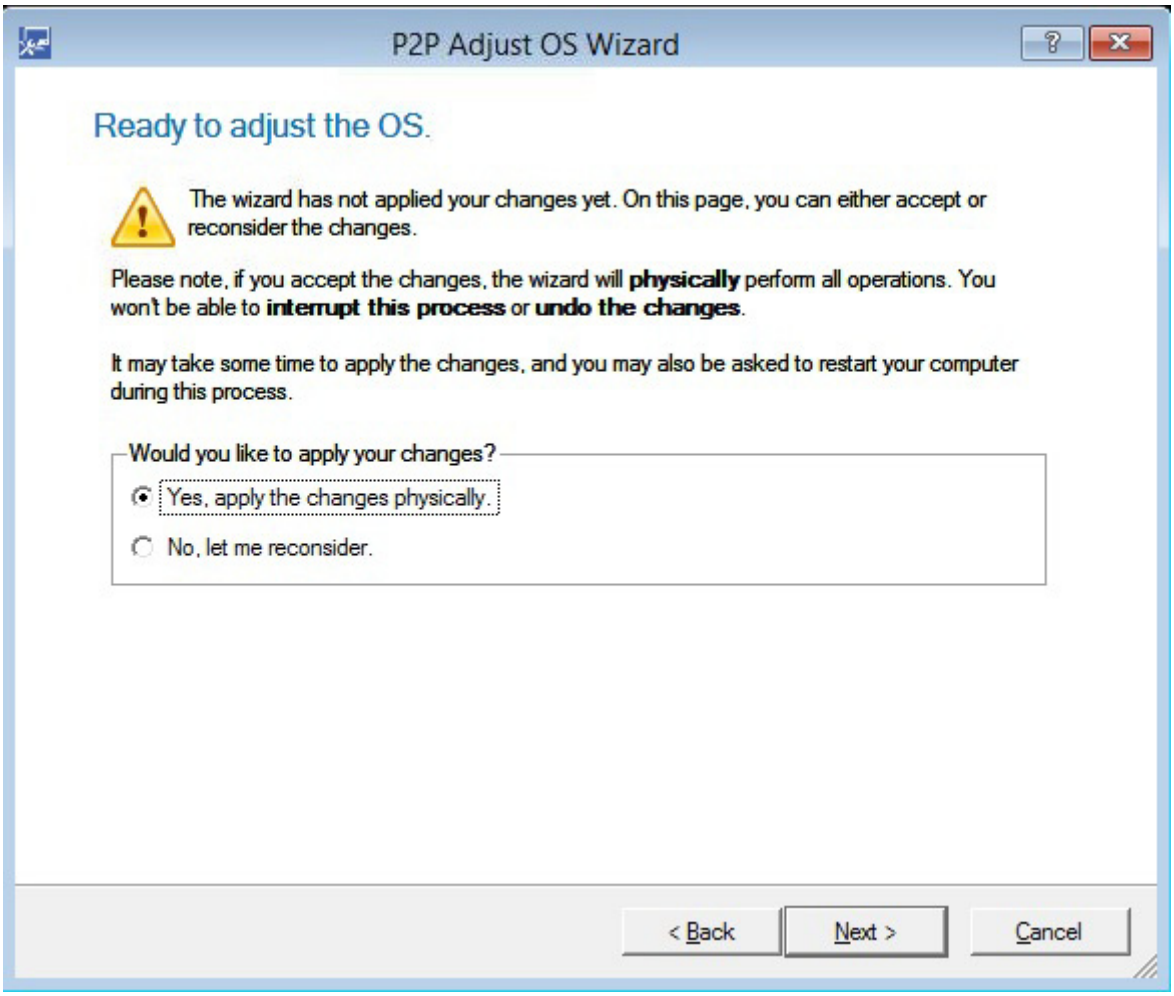
13. Do the following, after you verified that the restore completed successfully:

a. From the **P2P Adjust OS Wizard**, select the Recovery PN's operating system and, when prompted, choose **Adjust the OS to the new hardware automatically**, then **Next**.

If the Recovery PN's hardware is different from the image, the wizard needs to inject the necessary drivers into the image to adjust for this change.

Your snapshot might contain the necessary drivers in its cache; if so, this wizard will inject those drivers; otherwise, you'll need to supply the drivers.

b. When prompted, click **Yes apply the changes physically**, then **Next**:



c. Select **Reboot the system** or **Shutdown the system** to complete the restore, then **Finish**.

Troubleshooting

- If you reboot the PN without adjusting for new hardware, you can perform the adjustments after the restore completes. Simply launch the wizard again by repeating Step 3 and Step 4, then select the **Perform OS Adjustment** workflow from the **Continue**

onQ Restore page.



- If your PN will not boot, go to (Step 3) To manually adjust a PN's operating system for new hardware:.
- If your PN boots, but the onQ Portal indicates a **Connection Status** of `PN offline`, your PN might have acquired a DHCP IP address, not the original static IP address. Modify the PN's configuration to reflect its original, static IP address. Afterward, onQ can resume backups of the PN.

(Step 3) To manually adjust a PN's operating system for new hardware:

If after performing a BMR your PN does not boot (BSOD), you're missing boot-critical drivers. In the rare case that the P2P tool did not properly adjust your PN's operating system for the new hardware, you need to manually adjust the operating system using the P2P Adjust OS Wizard's **Set parameters for the OS adjustment** option. Before you begin, retrieve the necessary drivers.

1. Launch the QUARK wizard:

a. Configure the Recovery PN's BIOS to boot QUARK.

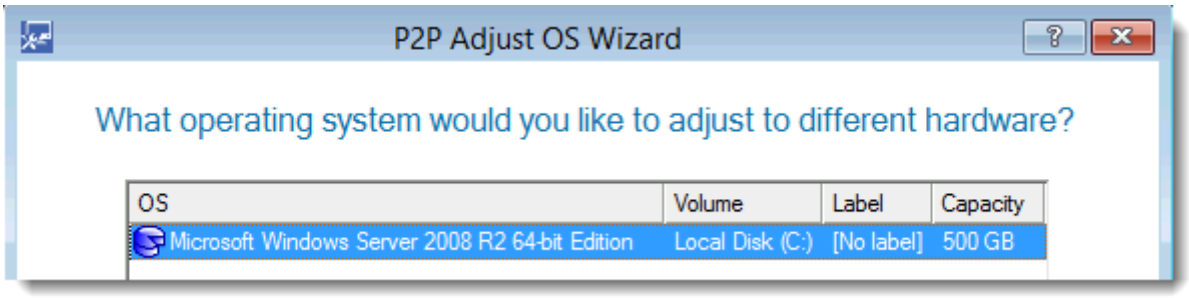
b. Boot the Recovery PN to load the QUARK image.

The **Start BMR** page appears.

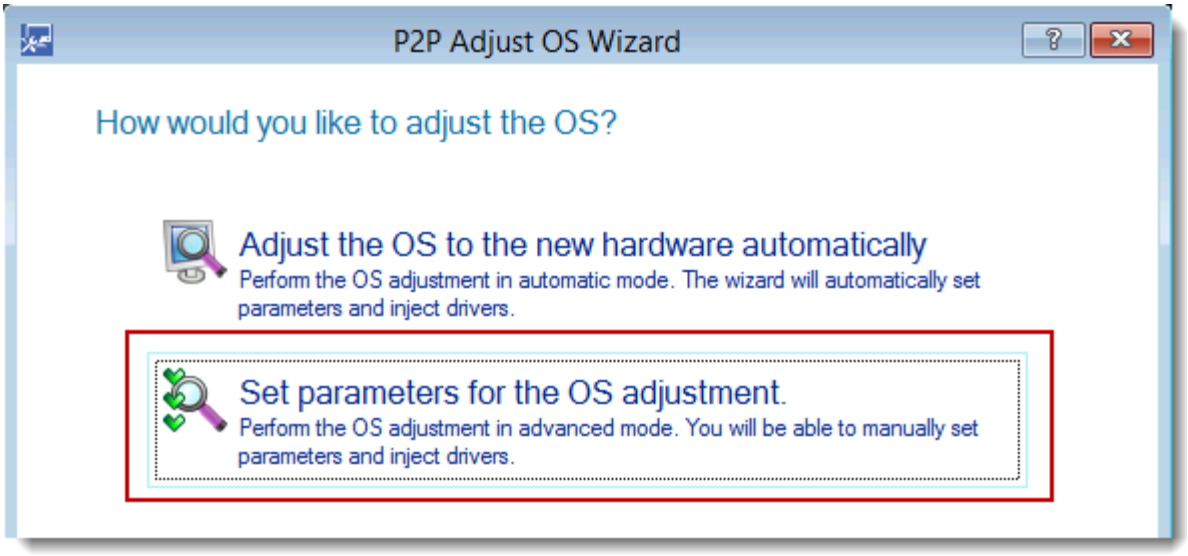
2. From the **Start BMR** page, allow the QUARK wizard to check for basic network and disk controller drivers on which QUARK depends, then **Next**.
3. From the **Continue onQ Restore** page, select the **Perform OS Adjustment** workflow.



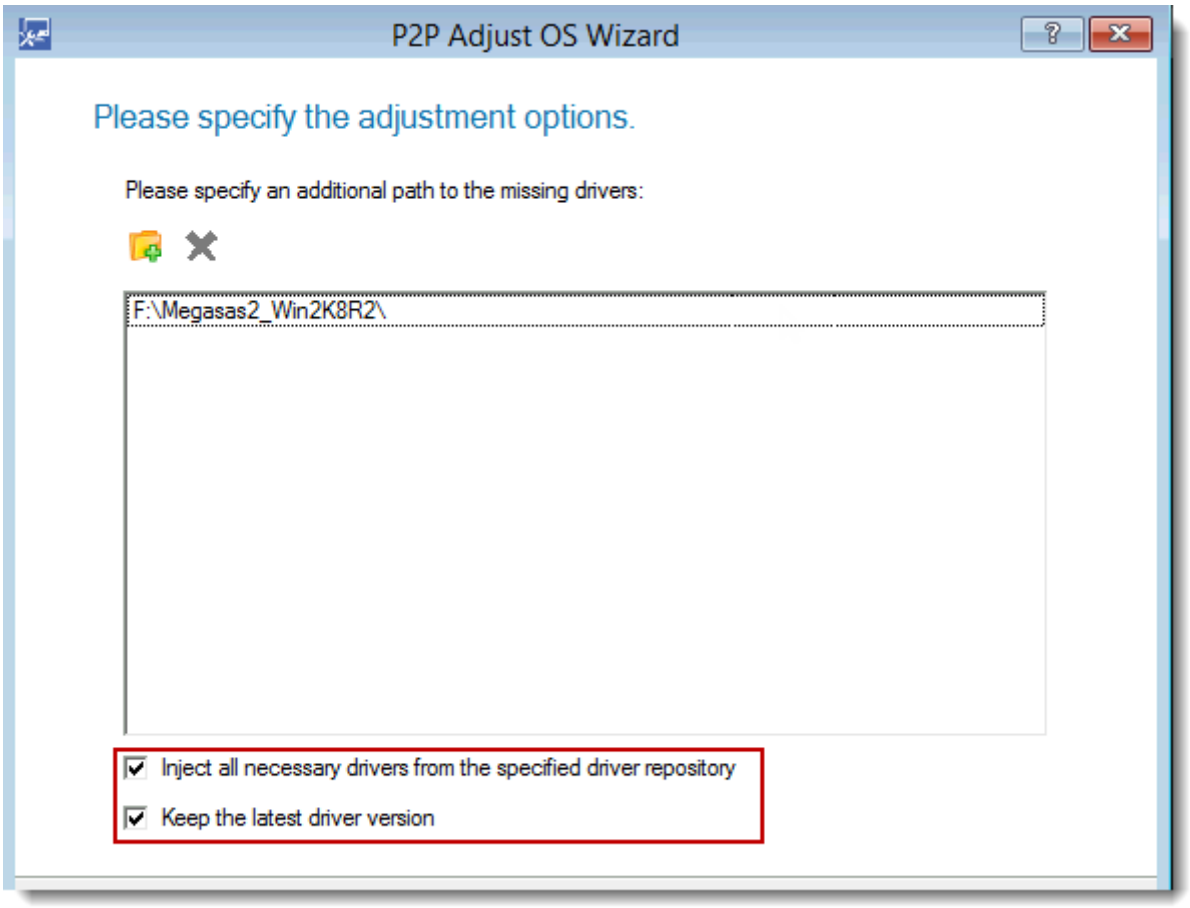
4. From the **P2P Adjust OS Wizard**, select the Recovery PN's operating system, then **Next**.



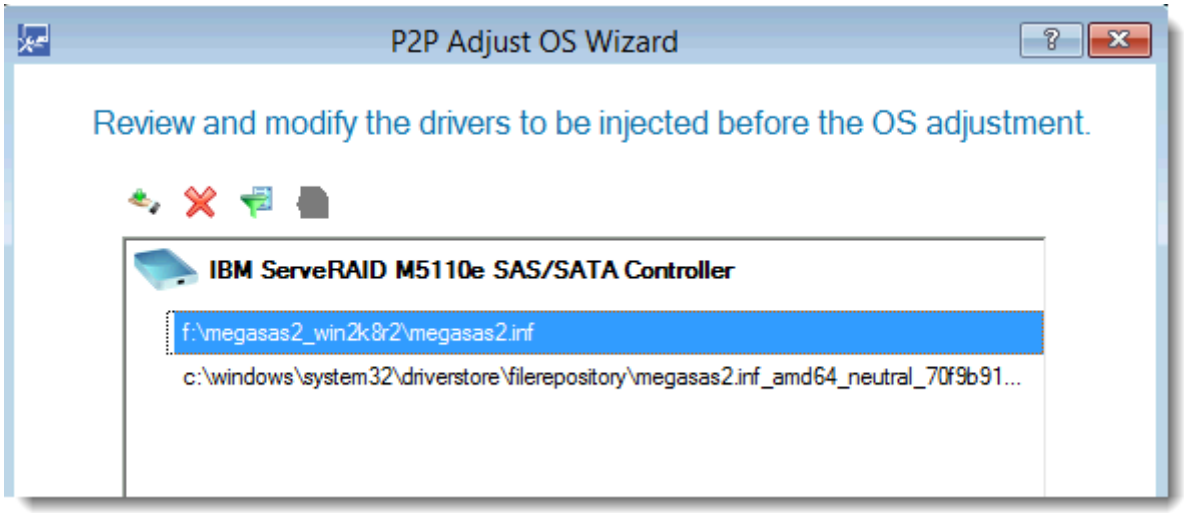
5. Choose **Set parameters for the OS adjustment**, then **Next**.



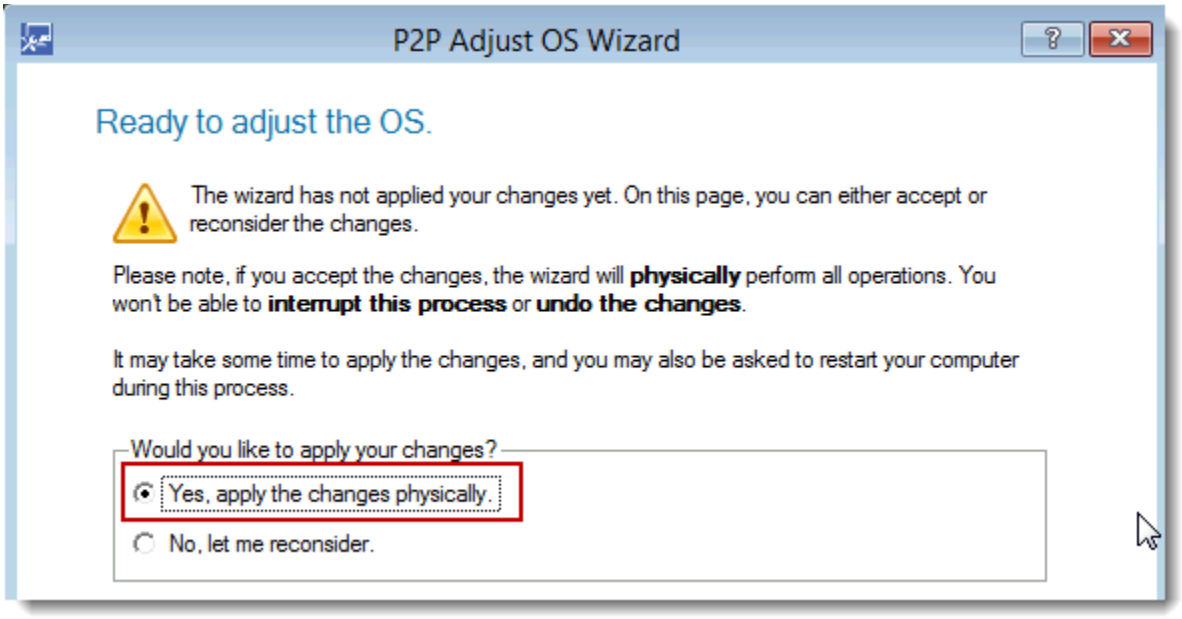
6. Specify the path to the driver, select the **Inject all necessary drivers from the specified repository** check box and the **Keep the latest driver version** check box, then **Next**.



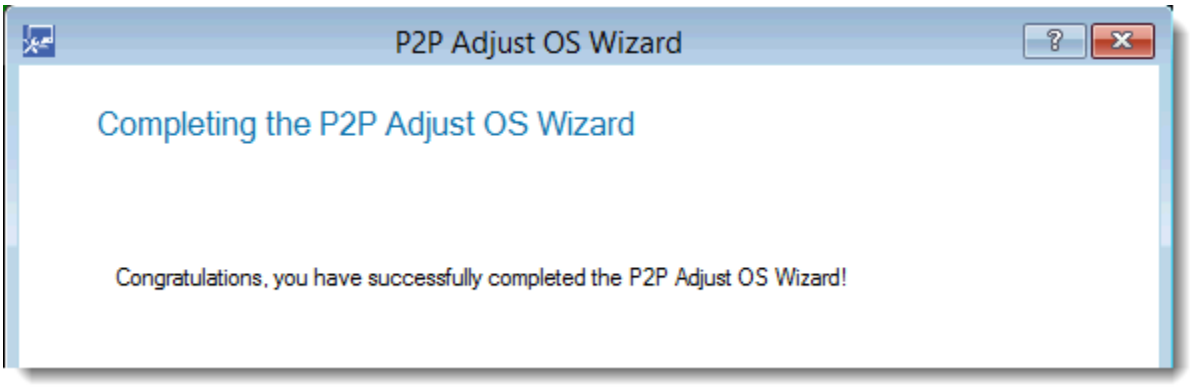
7. Select the driver, then **Next**.



8. Select the **Yes, apply the changes physically** radio button to make the necessary adjustments, then **Next**.



9. Click **Finish**. Now that your PN has the correct boot-critical drivers, your PN should now boot.



10. Select **Reboot the system** or **Shutdown the system** to complete the adjustments, then **Finish**.

(Step 4) To interpret QUARK Errors:

Error	Cause	What do I do?
`Authentication failed for onQ`	You typed the incorrect credentials.	Try again.
`Command Failed, IP Address is in use by a Protected Node`	When you specified an IP address for the QUARK client, you typed an IP address that's already listed in the onQ's hosts file. The QUARK client needs an IP address to communicate with onQ. A common mistake is to specify the PN's original IP address.	Change the IP address in QUARK's **Network Configuration** page.
`Failback already running for a different target`	The onQ Appliance is busy performing a BMR for a different Recovery PN. An onQ Appliance can only perform one BMR at a time.	Wait for the in-progress BMR to complete, then perform the next BMR.
`onQ Protection turned off. Could not continue with BMR process`	Protection on the onQ Appliance is off.	Turn on protection.

Error	Cause	What do I do?
onQ connection timeout	If the wizard doesn't get a response from the onQ Appliance within 60 seconds, the wizard displays this timeout error.	Try again.
Could not connect to onQ	The wizard could not connect to the onQ Appliance.	Try again.
Could not retrieve PN list from onQ	There was either an unexpected database error; or There are no PNs in the onQ Appliance's list of protected node (the list is empty).	Try again. If the problem persists, verify that you logged on to the correct onQ Appliance.
No Snapshots available for the selected PN	There are no snapshots available on this onQ Appliance for the PN you selected.	Verify that you logged on to the correct onQ Appliance.
onQ already started failback for a different target. onQ terminating current BMR session.	onQ allows concurrent BMR logins to the same onQ Appliance; however, precedence is given to the last BMR user. If you are not that user and you try to start a BMR/Incremental restore, you will be rejected by onQ: QUARK automatically redirects you to the logon page and displays the aforementioned message.	Perform the BMR at a later time.
EFI system partition not detected. Please check disk configuration.	BMR for EFI boot based PNs require EFI system partition to be created prior to restore. If the user's intent is to perform a EFI based restore, then QUARK checks the presence of EFI system partition before the start of the restore. QUARK returns this error if QUARK cannot detect an EFI partition.	Specify the correct partition style in the **QUARK Disk Configuration** wizard.

Related Topics

"Bare Metal Restore Support" in onQ Release Notes

Restore PNs Using QUARK

Perform File-level Restore

12

Monitoring

- Monitor multiple onQ Appliances
- Generate on-demand reports
- Monitor protected nodes
- Monitor recovery nodes
- Monitor backups
- Backup Activity Report
- Monitor DR Appliance
- Monitor onQ Archive Vault enrollment
- Monitor disk space and memory usage
- About Alerts
- Modify e-mail alert settings
- Retrieve logs

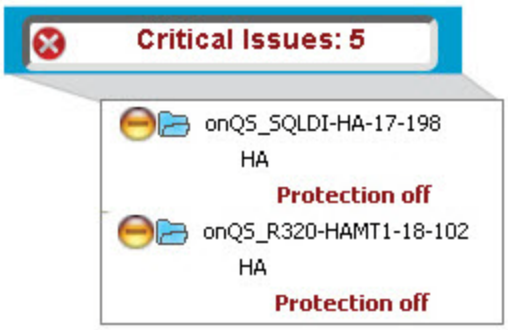
12.1 Monitor multiple onQ Appliances

Use onQ Monitor to monitor all of your onQ Appliances from a "single pane of glass". These onQ Appliances can be located in a local datacenter or in the cloud. Think of your onQ Monitor as a dashboard from which you can retrieve information about your onQ Appliances, Protected Notes, and Recovery Nodes.

onQ Monitor enables you to click one tab or one button to:

- retrieve status information and overhall health for all your onQ Appliances.
- display alerts about all your onQ Appliances.
- display current operations being performed on PNs and jobs.
- display repository use, RN disk in use, RN disk allocated for each onQ as well as RN disk usage/allocation for each PN, providing a quick overview of capacity utilization for each onQ Appliance.
- navigate to your onQ Appliance to retrieve more status information or to browse the log files directly.

For example, from one location you can view all the onQ Appliances that have protection turned off:



To log on to onQ Monitor.

1. Log on to the HA Appliance's onQ Portal.
2. Scroll to the bottom of the **DASHBOARD** page.

3. Click on the **onQ Monitor** link in the footer.



To configure onQ Monitor:

1. Log on to the onQ Monitor.
2. Add all additional HA Appliances and onQ Archive Vaults that exist in your environment:
   a. In the navigation pane, click **Add/Remove** button.
   b. Do one of the following to retrieve the the HA Appliance or onQ Archive Vault, then **Get Name**.
      - In the **HA/AV Address** text box, type the address of the HA Appliance or onQ Archive Vault
      - In the **Set Name** text box, type the address of the the HA Appliance or onQ Archive Vault.
   c. (Optional) Type a **Group** name if you'd like to assign the HA Appliance and its DR, DR Mirror(s), and the onQ Archive Vault to a group.

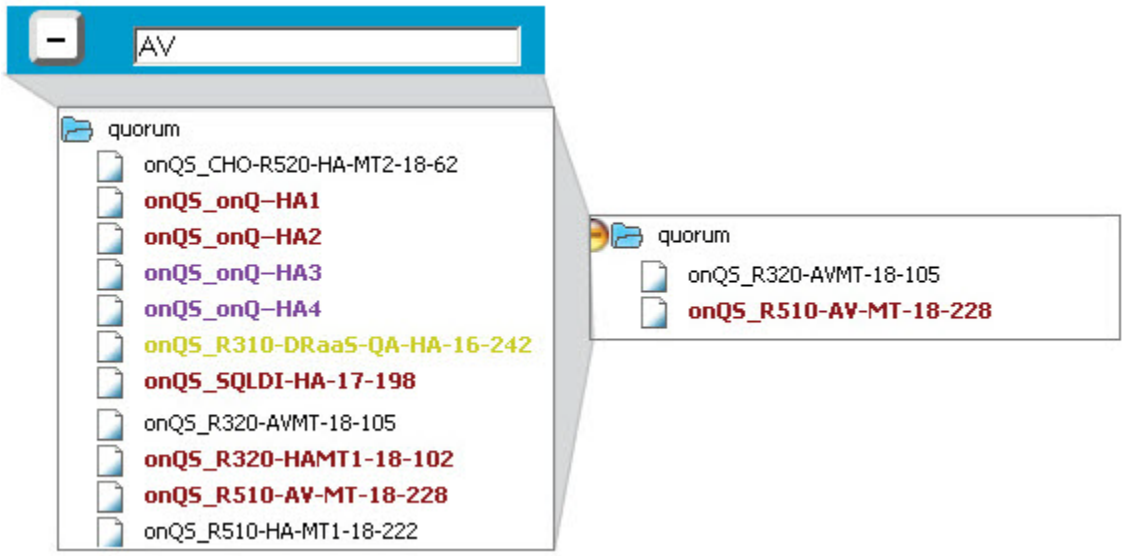
      Groups help you organize your onQ Appliances. For example, you can organize your onQ Appliances by domain, subnet, or department.

      By default, onQ Monitor provides a `quorum` group to include the HA Appliance from which you logged on to the onQ Monitor.
   d. Click **ADD**.

To filter on your onQ Appliances:

You can use filters to search for a specific appliance name, or if your onQ Appliances have a naming convention, you can filter on a string to search for multiple onQ Appliances. In the following example, a filter is performed on AV

string to search for all AV Appliances.



To delete a group:

When you delete a group, you also delete all the Appliance Sets (aka onQ Appliances) in that group.

1. Log on to onQ Monitor.
2. In the navigation pane, click **Add/Remove** button.
3. In the Group text box, type the name of the group (name is case sensitive), then **Delete Group**.

To browse the issues for a PN:

1. Log on to onQ Monitor.
2. in the navigation pane, click on the name of the onQ Appliance.
3. Click the **Issues** tab.
4. Wait while onQ Monitor requests issue information from the onQ Appliance. This retreival can take up to a minute depending on how busy the machine is and if you have a DR/DR mirror attached.

   onQ Monitor displays information for each PN by onQ Appliance. Each onQ Appliance is a hyperlink.

To browse the status of a PN:

1. Log on to onQ Monitor.

2. in the navigation pane, click on the name of the onQ Appliance.
3. Click the **Status** tab.
4. Wait while onQ Monitor requests status information from the onQ Appliance. This retreival can take up to a minute depending on how busy the machine is and if you have a DR/DR mirror attached.

   onQ Monitor displays information for each PN by onQ Appliance. Each onQ Appliance is a hyperlink.

To browse the status of a job:

1. Log on to onQ Monitor.
2. From the navigation pane, apply a filter to display all AV Appliances. Go to To filter on your onQ Appliances:
3. Click on the **STATUS** tab.
4. Wait a few seconds while onQ Monitor requests status information from the AV Appliance.

   onQ Monitor displays information for each job by onQ Appliance.

Related Topics

Generate on-demand reports

12.2 Generate on-demand reports

onQ Monitor enables you to generate reports that provide you both realtime and historical data. onQ Monitor aggregates this data across one or more onQ Appliances, highlighting the current activities on your onQ Appliances and PNs and the trends over the last 2 years.

These reports (charts and graphs) help you answer questions about your system in an effort to help you:

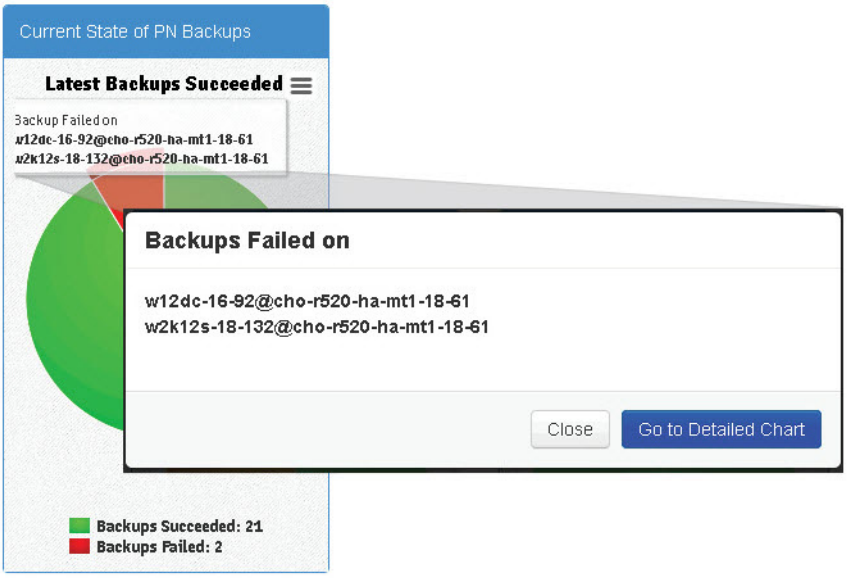
- make changes to your systems.
- forecast the need for system changes.
- prevent system failures.

• debug system issues.

To browse an Overview Report:

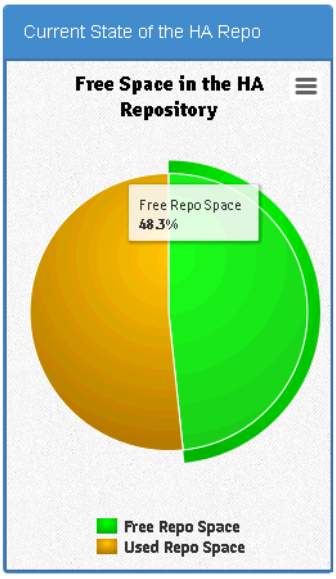
Overview Reports represents real-time data. Overview Reports aggregate the most recent backup, repository information, transfer, and self test for all the PNs. Unlike Historical Reports, Overview Reports do not display historical data; however, you can double-click on the pie charts to skip to the related historical report.

- **Backups Succeeded Chart**. This chart displays whether the *most recent* (last) backup for every enrolled PN was successful or not. If, for example, 21 backups out of 23 were successful, the success rate is 91%. Simply hover-over or double-click on the pie chart to retrieve a list of the PNs that correspond to the backup failures.



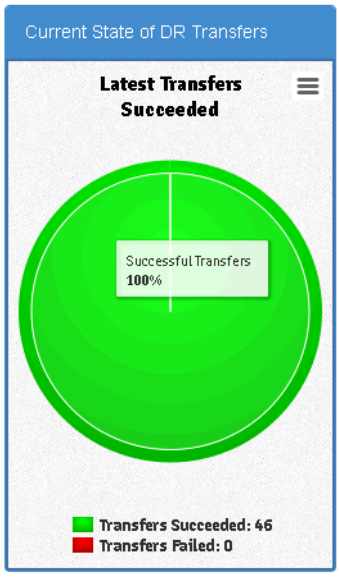
- **Free Space in Repository Chart**. This chart displays the *current* percentage of both free and used disk space in the snapshot repository. In case of a multi-tenant HA Appliance, the chart displays the sum of the free space on each snapshot repository and the sum of the used space on each snapshot repository Hover-over text provides you the percentage of free and unused

space.

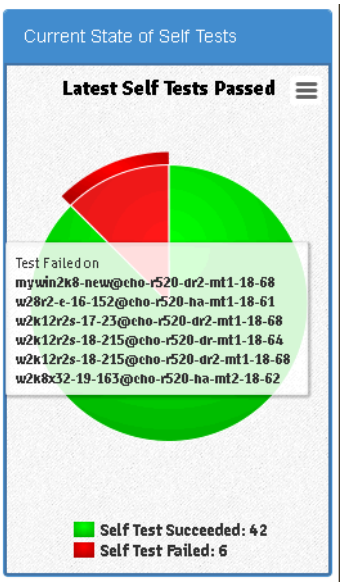


- **DR Transfers Succeeded Chart**. This chart displays whether the *most recent* (latest) DR transfer for every enrolled PN was successful or delayed. If, for example, you have 4 enrolled PNs, the chart displays the aggregation of the latest DR transfer for each of those PNs. If the DR transfer was successful for one of those PNs, but the DR connection was severed for some reason, the DR transfer for the three other PNs will be delayed until the connection is operating; in this case, this chart displays 75%

delayed.

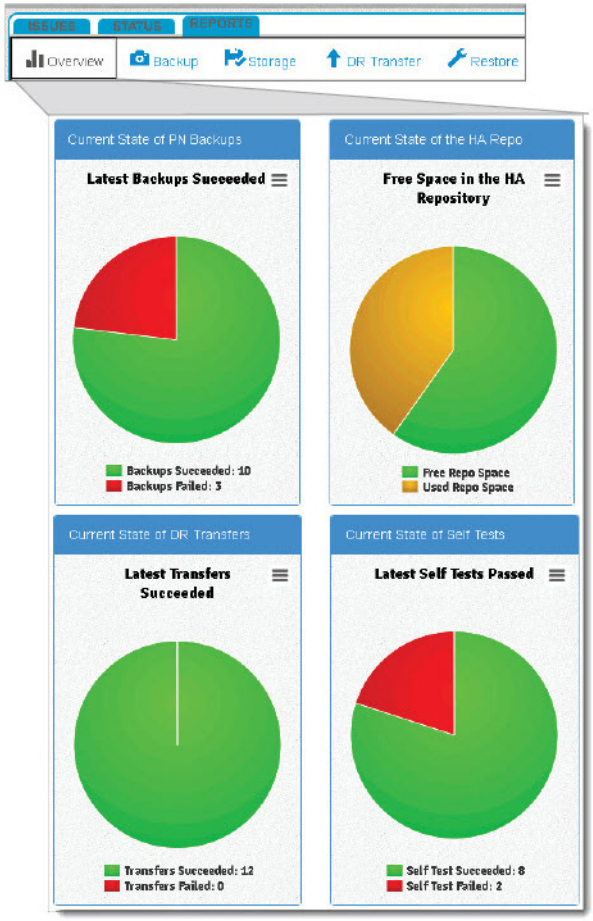


- **PNs Passing Self Test Chart**. This chart displays the *most recent* (latest) self test that passed for every enrolled PN. If, for example, you have 42 enrolled PNs and 6 of them failed the latest self test, the pass rate is 88%. Simply hover-over or double-click on the pie chart to retrieve a list of the PNs that correspond to the self test failures.



1. Log on to onQ Monitor.

2. Click on the **REPORTS** tab > **Overview** subtab.



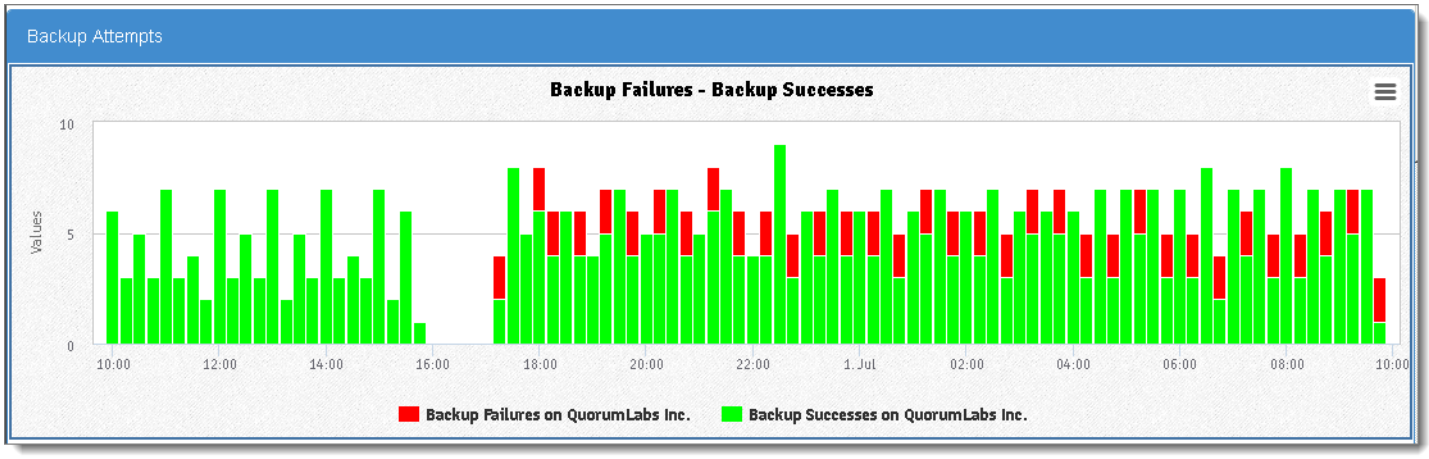
3. Wait for the charts to display.

To generate a Historical Report:

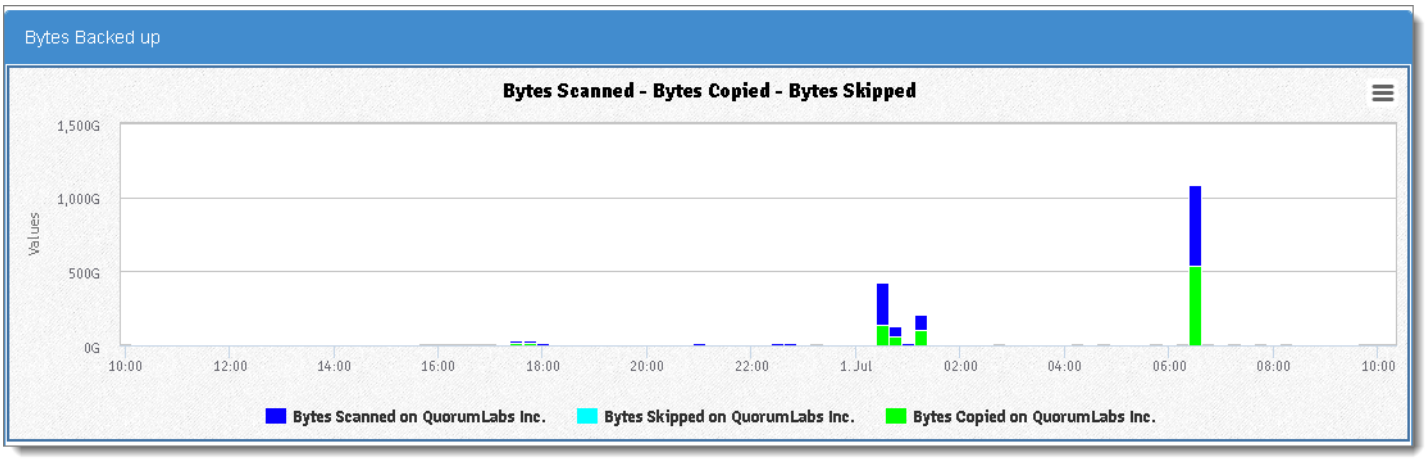
Historical Reports can report on activity going back 24 months. Historical Reports communicate reporting information using graphs. You can choose from the following Historical Reports:

- **Backup Historical Report**. Use this report in conjunction with the knowledge that queued snapshots can occur if the onQ is overly busy backing up other PNs or if the PNs have short backup intervals.
    - **Backup Attempts Graph**. This graph shows the total number of backup failures and successes in the time interval that you

specified across different instances.

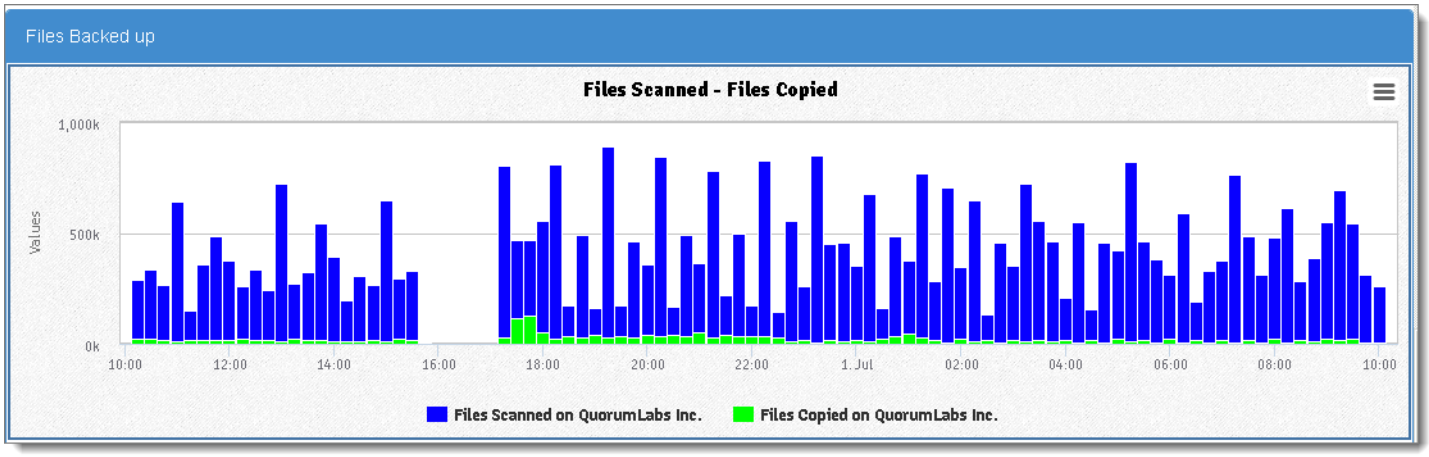


- **Bytes Backed up Graph**. The total number of bytes squirtcopy scanned, copied, and skipped in the time interval that you specified across the different instances. onQ skips bytes/files if instructed to do so as outlined in Edit backup exclude list.

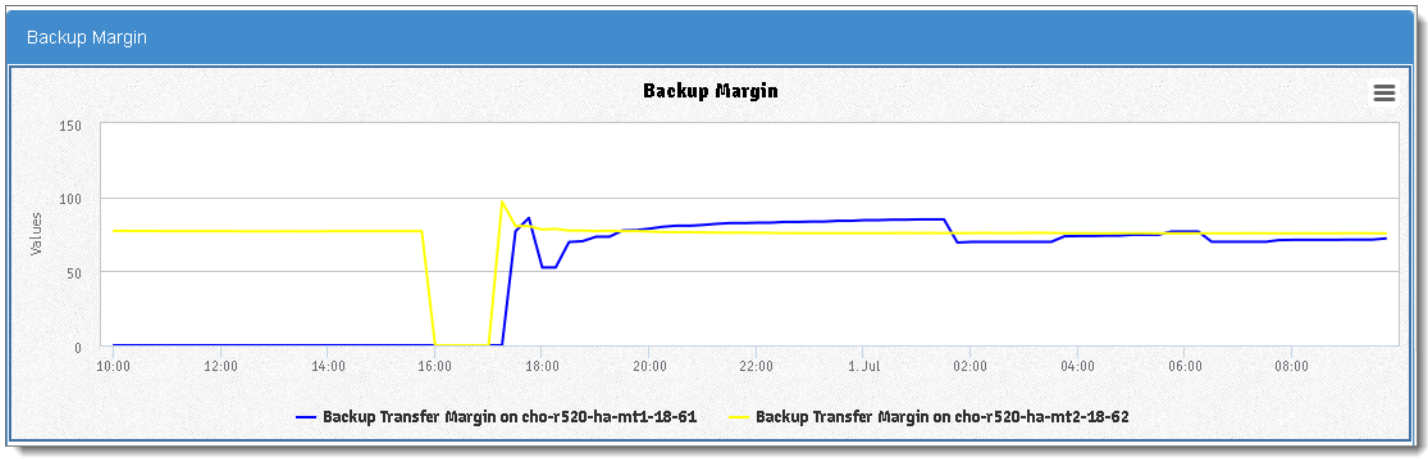


- **Files Backed up Graph**. The total number of scanned files on the PN by squirtcopy and the total number files that squirtcopy copied

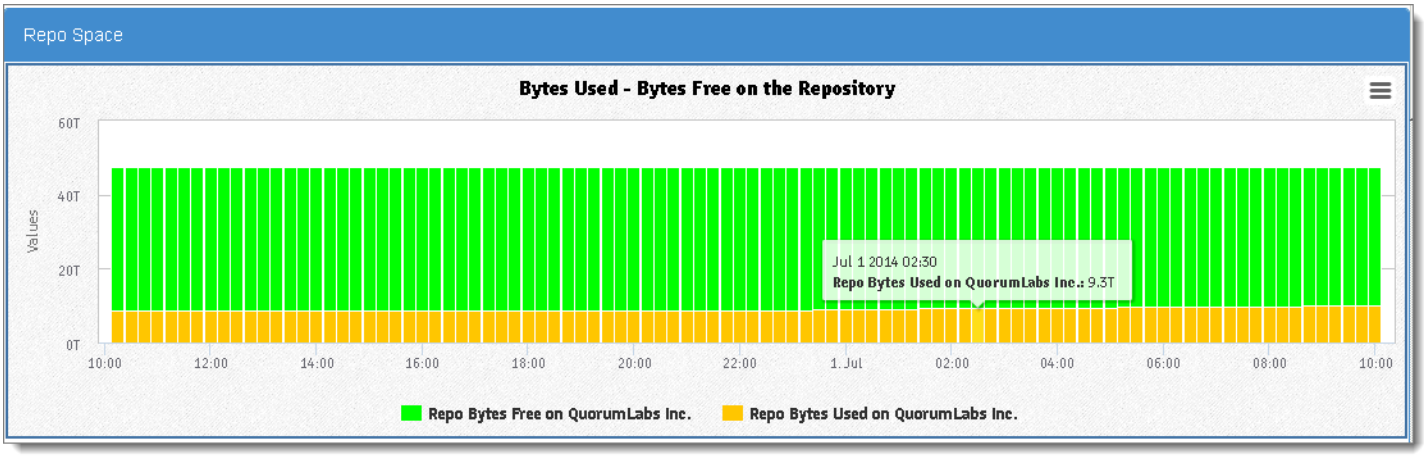
to the instances and within the time interval that you specified.



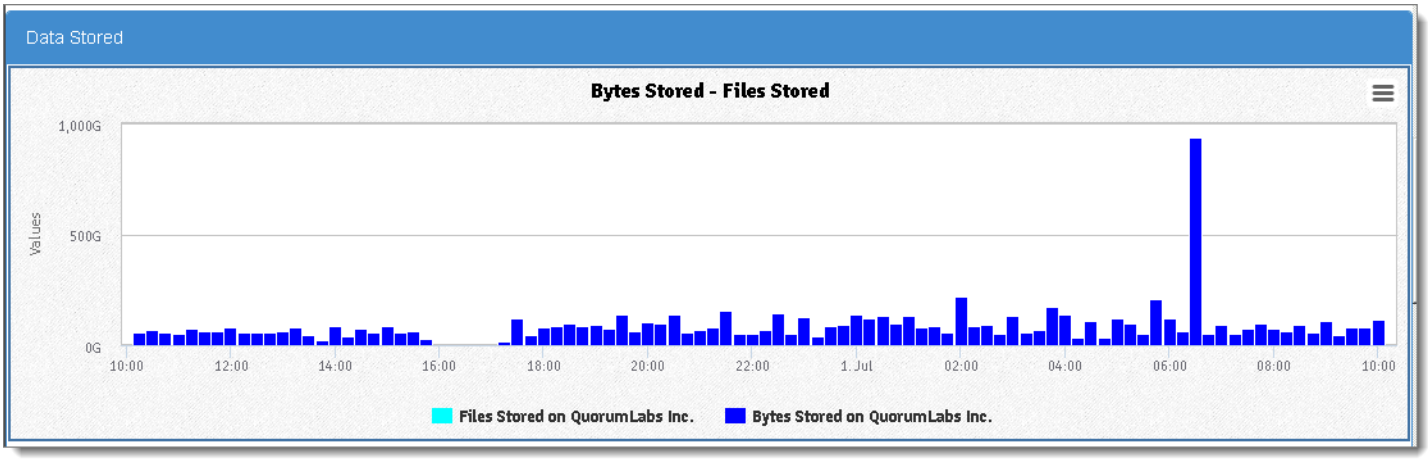
- **Backup Margin Graph**. If the margin value is low, backups are taking longer than expected to complete. This is the same value shown in the HA portal's **Backup Transfer Margin** column. You might need to change your backup interval as outlined in Monitor protected nodes and Monitor backups.



- **Storage Historical Report**. Use this historical report to determine if you need to add more snapshot disk space or delete orphan backups manually to free up space.
    - **Repo Space Graph**. This graph displays the total number of bytes used on the Repository.

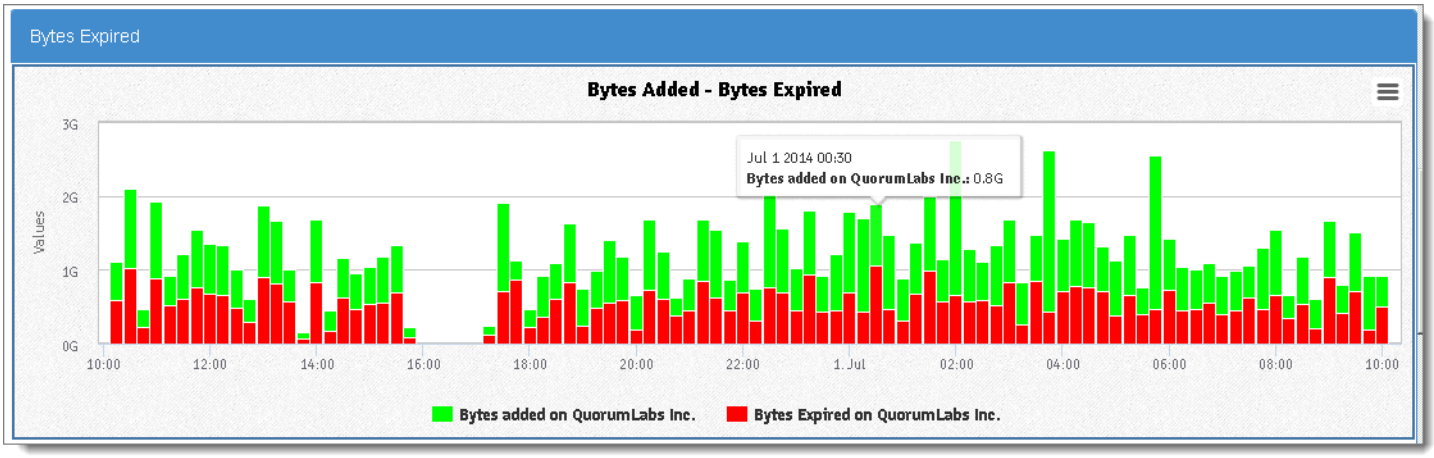


- **Data Stored Graph**. This graph displays the total bytes or files stored in the snapshot repository. The repository is the deduplicated archive of all snapshots of all of your Protected Nodes. From the repository, onQ can reproduce a complete system image from any snapshot.

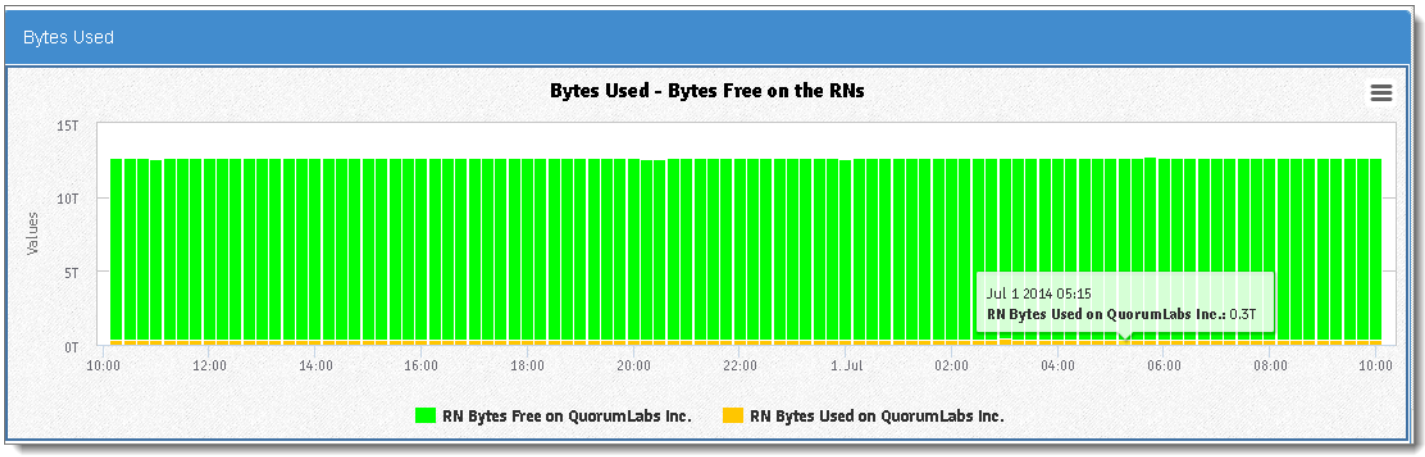


- **Bytes Expired**. When the onQ Appliance was pre-configured, a specific amount of disk space was allocated to the repository. Moreover, onQ has a built-in policy whereby if its disk space exceeds 85% onQ begins to "free up" space by deleting PN backups, starting with the oldest backups. Files expired and bytes

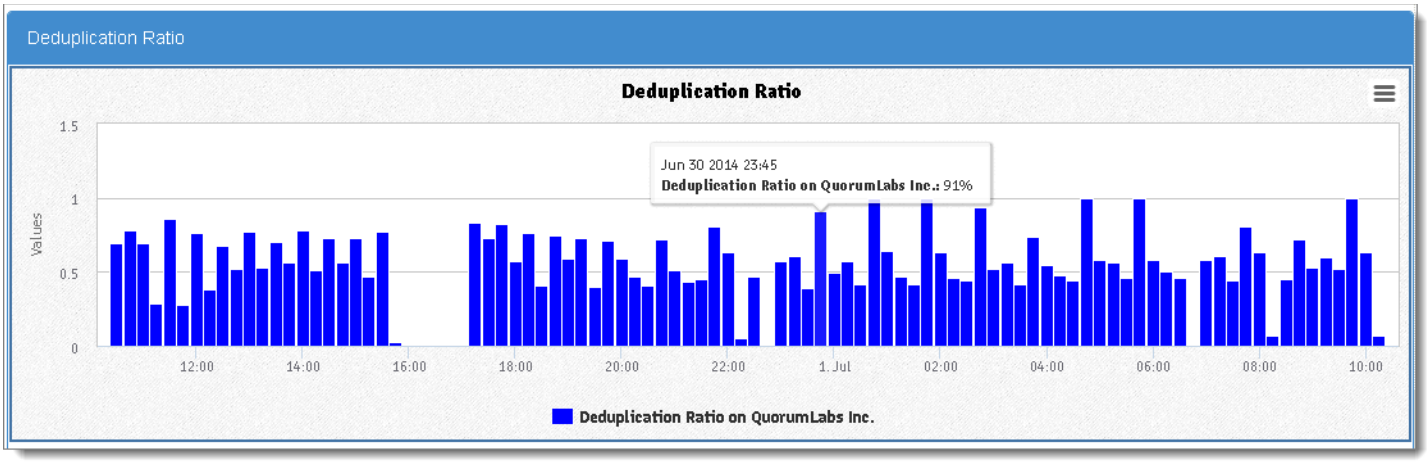
expired values refer to the data (in files and bytes) that onQ expired plus the bytes expired due to the retention period being met.



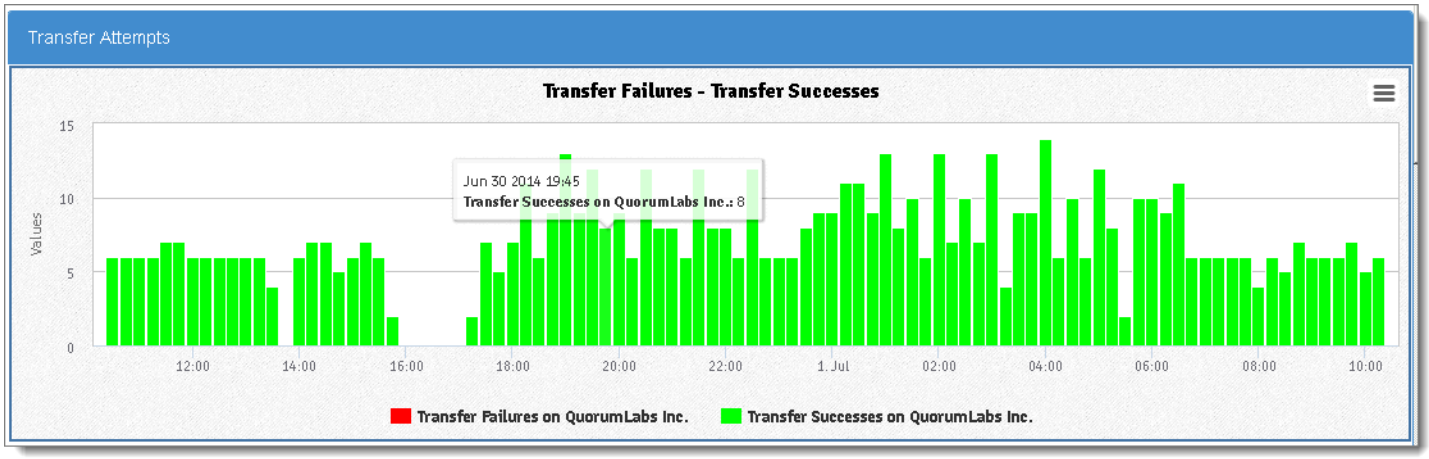
- **Bytes Used Graph**. This graph displays the total number of bytes used on the Recovery Nodes.



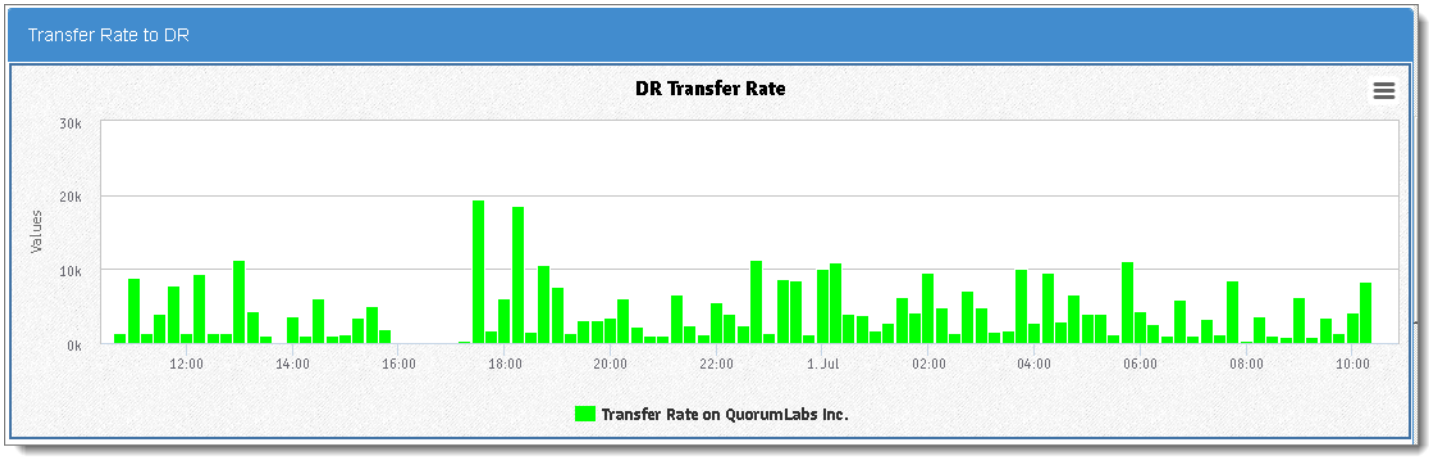
- **Deduplication Ratio Graph**. This graph shows the data compression ratio, which is equal to the total number of bytes copied divided by the total number of bytes added. If your PNs all have the same data, you'll see a better compression ratio then if your PNs have unique data; the latter scenario is indicative of duplicate data on your PNs.


Deduplication Ratio
Deduplication Ratio
Jun 30 2014 23:45
Deduplication Ratio on QuorumLabs Inc.: 91%
Values
1.5
1
0.5
0
12:00
14:00
16:00
18:00
20:00
22:00
1. Jul
02:00
04:00
06:00
08:00
10:00
Deduplication Ratio on QuorumLabs Inc.

- **DR Transfer Historical Report**. After performing a backup, the HA Appliance transfers this data from the HA Appliance to the DR Appliance. Use this report to help determine whether or not to adjust bandwidth throttling to speed up data transfers.
    - **Transfer Attempts Graph**. This graph shows the number of HA-to-DR or DR-to-DRMirror or both transfer attempts and successes in the time interval that you specified and depending on the instances that you choose. For example, if you want to show all HA-to-DR transfers, then select the HA Appliances; do not select any DR Appliances.

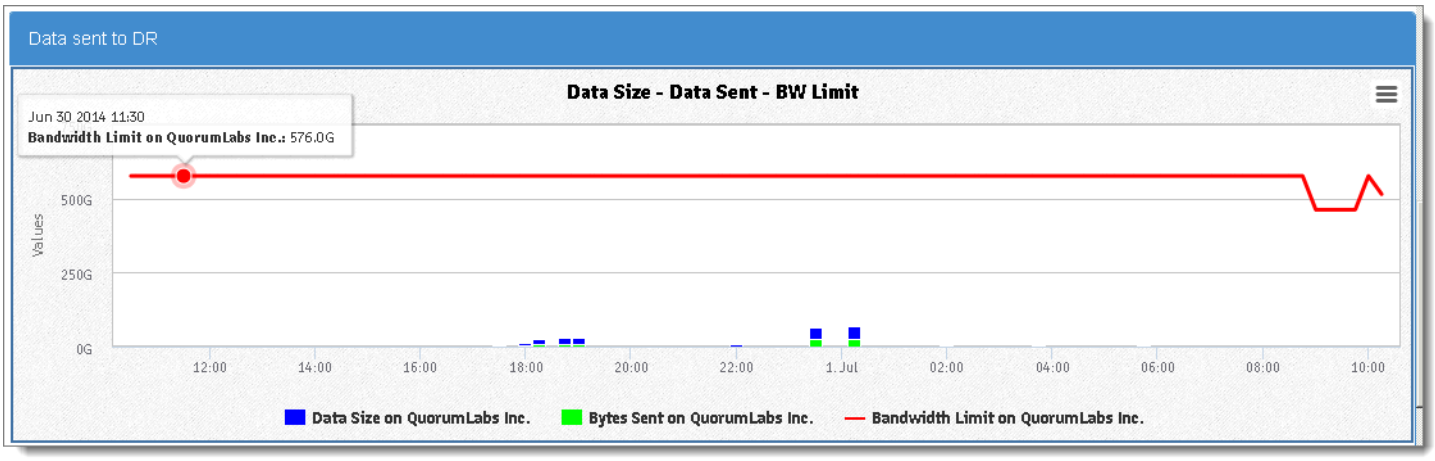


    - **Transfer Rate to DR Graph**. This graph shows the transfer rate, in bytes/seconds, for all successful transfers within a time interval that you specified.

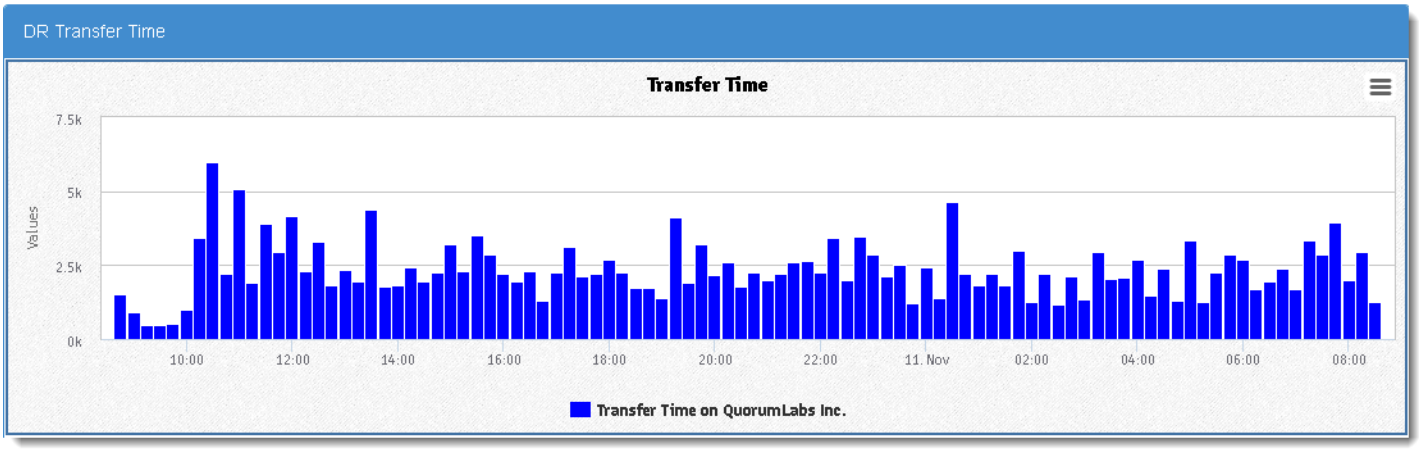


    - **Data Sent to DR Graph**. This graph shows the total bytes sent from the HA appliance and the total bytes sent to the DR

appliance. The bandwidth limit is the transfer limit at the time of transfer that you defined in the onQ portal.

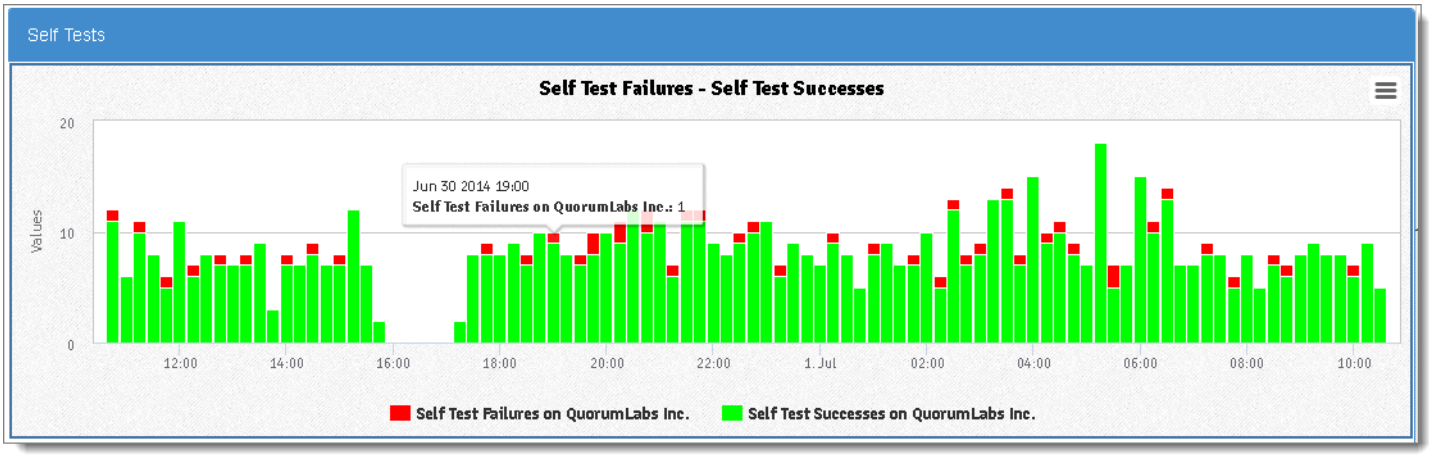


- **Transfer Time**. If the transfers are taking longer than expected to complete, you might need to modify the backup interval, modify the Window Start time, and adjust bandwidth throttling.

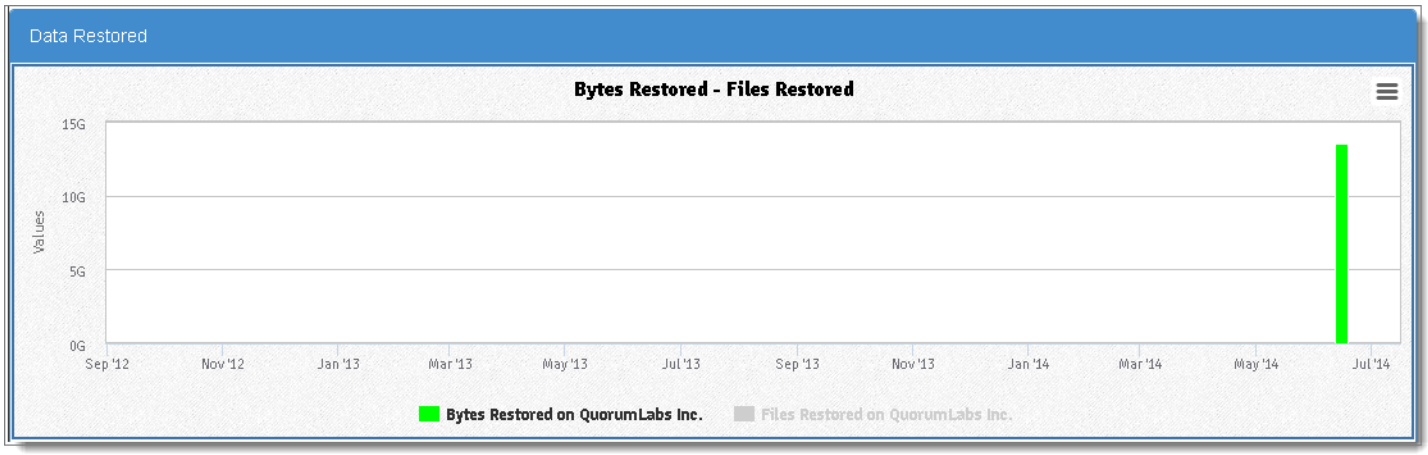


- **Restore Historical Report**.
  - **Self Tests Graph**. onQ can automatically test (also called Self Test), on a predefined schedule, recovery nodes' ability to boot properly and can communicate with the onQ Service. This historical report shows the number of self test failures and successes in the time interval that you specified. Use this data to

help assess your PN's vulnerability to disasters.



- **Data Restored Graph**. This historical report shows the number of bytes and files restored using FLR. You can use this chart to provide restore statistics for SLA reporting. You can also use this report to verify that the number of files and bytes that you expected to restore matches the files and bytes that were actually restored.

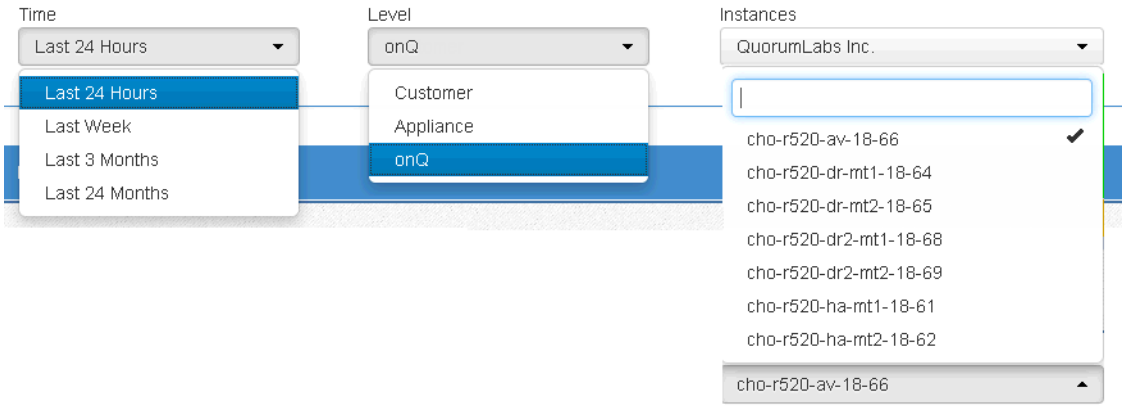


- **Restore Attempts Graph**. This historical report shows the

number of FLR restore attempts that failed and succeeded.



1. Log on to onQ Monitor.
2. Click on the **REPORTS** tab.
3. Click on the **Backup**, **Storage**, **DR Transfer**, or **Restore** subtab. Each of these pages maps to a different Historical Report.
4. Specify the reporting parameters.



a. In the **Time** drop-down list, specify the time interval that you want to capture. Each time interval displays information in specific reporting increments.

Time Interval	Reporting Increment
Last 24 Hours	15 minutes
Last Week	2 hours
Last 3 Months	Daily
Last 24 Months	Weekly

b. In the **Level** drop-down list, specify the entity on which you want to report:

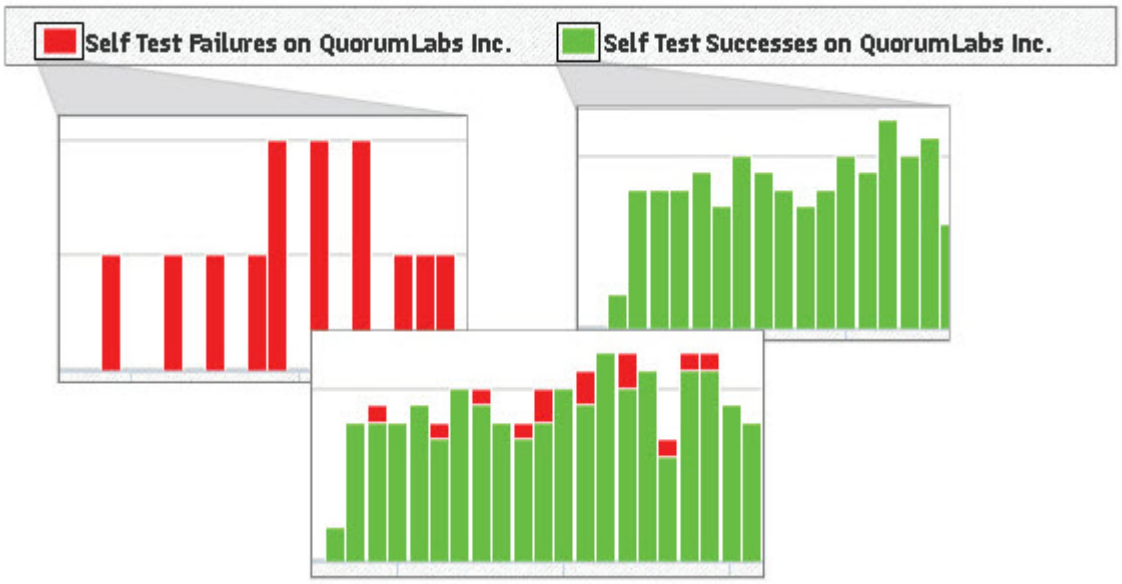
Level	Displays data for...
Customer	All of a customer's onQ Appliances combined. For example, HA and DR. Customer is defined by **APPLIANCE CONFIG** > **ADVANCED** > **License** > **Customer** field.
Appliance	One specific onQ Appliance. The onQ Appliance is identified by **APPLIANCE CONFIG** > **HYPERVISOR** > **Hypervisor Name** field.
onQ	One specific onQ. In the case of a multi-tenant environment (**DASHBOARD** tab > **onQ STATUS** page > **Configuration** field), one onQ Appliance can have many onQs. The onQ is identified by the onQ Manager hostname (aka Manager hostname or onQ hostname) as shown in **DASHBOARD** tab > **onQ STATUS** page > **Manager hostname** field.
PN	One specific protected node.

c. In the **Instances** drop-down list, select the specific Customer(s), Appliance(s), onQ(s), or PN(s) that you want included in your report.

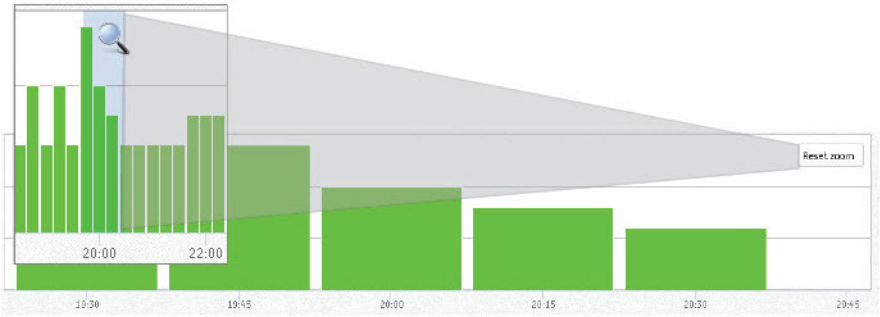
5. Click **Update Chart**.

6. (Optional) Click on the Legend links within the report to include or exclude specific data (quantities). You can use this feature to compare data.

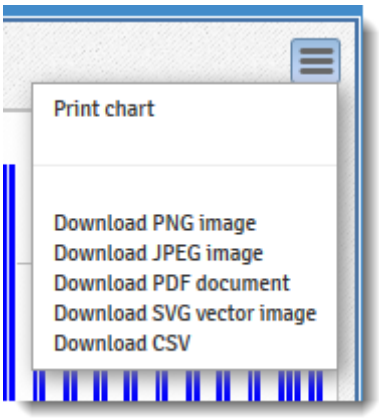
In some cases, you'll need to exclude data (toggle between legend links) when there is a huge disparity in the quantities. For example, you cannot view the files restored and bytes restored simultaneously in the Restored Data Chart because the scale for number of restored bytes has a long range relative to the number of files restored thereby hiding the bar color for the files restored.



7. (Optional) Zoom in on a subset of data by highlighting a section of the timeline. Click the **Reset Zoom** button when you're done to return to the entire report.



8. (Optional) Export the report. An export does not include zoom-in data.



12.3 Monitor protected nodes

From the onQ Portal you can quickly identify the connection status, protection status, and backup status of protected nodes.

Monitoring is particularly useful if you've started a protected node or you recently upgraded a host.

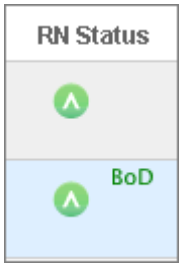
To monitor the status of a protected node:

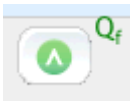
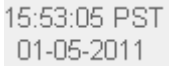
1. Log on to the HA Appliance's onQ Portal.
2. Go to **DASHBOARD** tab > **PROTECTED NODES** page.

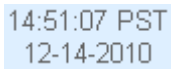
   A list of all your protected nodes and their current status appears:

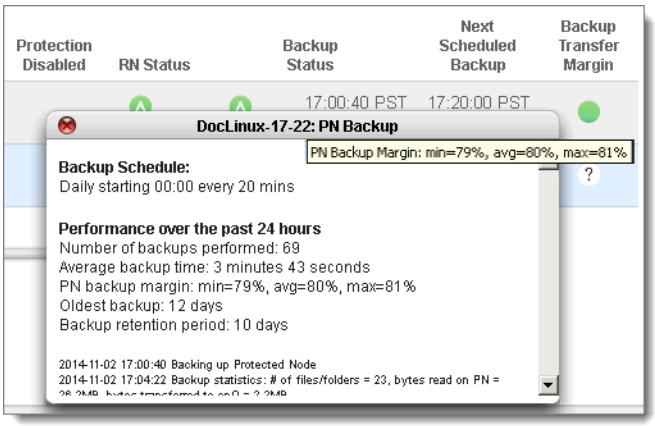
Connection Status	 carat is green and white arrow points upward	PN is running and protection is turned on. Tooltip indicates `PN online` (agent-based) or `Proxy online` (agent-less). • (agent-based) the online state of the PN reflects the Power State of the PN. • (agent-less) the online state of the Proxy reflects the communication state of the PN proxy on the ESX Server.
	 carat is green with a yellow dash	For agent-less PNs, the PN proxy is currently being upgraded: 1. Before upgrade  : solid green carat; white arrow points upward. 2. During upgrade  : solid green carat; yellow dash. 3. After upgrade  : solid green carat; white arrow points upward.
	 carat is red and white arrow points downward	PN is not online. Tooltip indicates `PN offline` or `Proxy offline`. See (Agent-based PNs) Connection Problems or (Agent-less PNs) Connection and Backup Problems.
		onQ Appliance's connection to PN is secure.
		onQ Appliance's connection to PN is not secure. See "(Agent-based PNs) Connection Problems" on page 389.
Protected Node		Hostname of the protected node. Hostname is the same as the Windows hostname.
Type	**PN**	PN itself is running as opposed to the recovery node.
	RN 	RN, not the protected node is running, implying that the primary node has failed or has been stopped and its recovery node equivalent has been started.

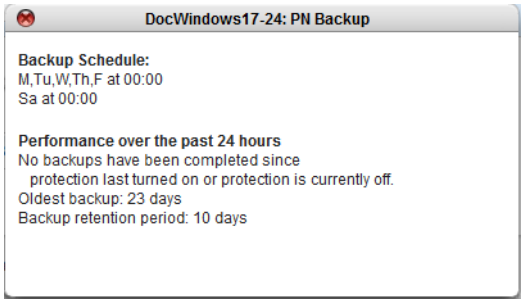
Protection Disabled	no check mark	Protection is enabled on this node.
	✓ check mark appears	Protection is disabled on this node.

RN Status	 carat is green and white arrow points upward	Recovery Node image of the Protected Node is available and ready to run. The acronym `BoD` appears next to the icon if `Auto RN Creation` (see Modify recovery node creation policy) is set to disabled.  Mouse-over help displays a record of recent events for the RN.
	 carat is green and white arrow moves top to bottom	onQ is preparing the Recovery Node. Mouse-over text indicates. `RN is not updated – RN is NOT Ready to Run`.
	 carat is red and white arrow points downward	No Recovery Node is available. See Recovery Node Problems. The acronym `BoD` appears next to the icon if `Auto RN Creation` (see Modify recovery node creation policy) is set to disabled.
	 hourglass on a yellow background	Internal processing is occurring. Typically this icon persists for one to two minutes before changing to red or green.
	 carat is green, white arrow points upward, and is accompanied by a red asterisk	The Recovery Node is not running the most recent snapshot. The RN is more than one snapshot behind the most recent snapshot in the Repository. Mouse-over text displays the snapshot that the RN is running. The red asterisk might appear in a normal backup cycle for a short period of time, after the new snapshot is created and before the update of the RN begins.

Backup Status	 carat is green and white arrow points upward	Complete restorable backup exists in the repository. Mouse-over text indicates `Idle, backup complete`. The date and time of the most recent completed backup appear to the right.
		The Qf icon indicates that the filter driver is enabled.
		The Qf icon indicates that the filter driver is enabled, but the red asterisk at the end of the Last Backup time indicates that onQ did not use the filter driver when performing the last successful backup. Due to problems with the logs, onQ resorts to using the Non-Filter backup method to correct the problem. The PN's Event Logs also confirm this proactive action. onQ uses Qfilter for the next backup and the red asterisk disappears. The Event Log shows the correct Qfilter state in the backup statistics.
	 carat is green and white arrow moves left to right	Backup is in progress. Mouse-over text indicates `Running` or Backup
	 carat is green and white arrow moves top to bottom	Transfer of the backup to the repository is in progress. Mouse-over text indicates `Adding backup to repository`.
	 carat is green, white arrow points upward, and accompanied by a red asterisk	A valid snapshot exists, but the last backup attempt resulted in an error or the snapshot wasn't added to the repository.
	 carat is red and white arrow points downward	Not a single copy of a backup snapshot exists.
	 timestamp	Date and time of the most recent successful backup; this timestamp is equivalent to the corresponding log entry: Backup time end = "Backup complete."

Next Scheduled Backup	 timestamp	Date and time of the next scheduled backup. Matches the output in the corresponding log entry: Backup next schedule =”nxt_<pnName>:<timeStamp>:<timeStamp>”.

Backup Transfer Margin	solid green icon	Backup margin is within an acceptable range. Backup margin indicates how well PN backups are keeping up with the backup interval as calculated by averaging the minimum, average, and maximum intervals from the previous 24 hours. This 24-hr time period automatically resets each time you restart protection. Backup margin is the amount of free time between the completion of one backup and the start of the next. Backup margin uses three inputs from the Event Log to calculate the minimum (min), average (avg), and maximum (max) margins that display when you mouse-over the Backup Margin icon. • `Backing up Protected Node` = Backup start time • `Backup complete` = Backup end time • `Next Backup scheduled for <timeStamp>` = Next Scheduled Backup In other words: `Backup margin = (100 - (100 * ((endTime-startTime)/nextScheduledBackup-startTime)))` So as the backup takes longer, the backup margin will decrease. Incremental backups take longer when the PN is performing continuous database updates. Mouse-over the icons to display more information:  For aggregate data on backup transfers, see the DR Transfer page header.

Backup Transfer Margin (Contd)	solid yellow icon	Backups transfers are not within an ideal range—the backup margin *average* is below 35%. Modify the backup interval.
	solid red icon	Backups transfers are not within an acceptable range—the backup margin *average* is below 10%. Modify the backup interval.
	solid grey icon	The PN was configured with replication disabled. Mouse-over help indicates `DR Transfer Margin: N/A` and `DR Replication Disabled`. Tooltip dialog's title read `Replication Disabled`.
	white icon with question mark	Backup margin is not applicable because no backups have been performed in the last 24 hours. Also displays if protection is off or if backup or DR transfer has not been completed since you turned on protection. 
Disk Usage	3 %	Percentage disk usage of entire onQ Appliance, and includes any recovery nodes that are running.

Related Topics

Monitor recovery nodes
Monitor backups
Generate on-demand reports

12.4 Monitor recovery nodes

Each of your primary protected nodes has a corresponding recovery node. The HA Appliance constructs a recovery node after each periodic backup.

The recovery node is then ready to be started as a replacement for the

primary if the protected node fails. Even after you have started a recovery node, backups continue on the HA Appliance.

If the RN has been started and is running in production mode, it effectively becomes a PN; the RN, not the PN, is being backed up by the HA Appliance.

In the event that a protected node fails, you need to start the corresponding HA recovery node. Similarly, if there's an HA recovery node failure or HA Appliance failure, you need to start the recovery node on the DR Appliance. In both cases, you'll want to monitor the progress of these recovery nodes. Of course, in the event of a disaster, you'll receive alerts.

To view the status of a recovery node:

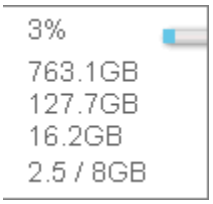
1. Log on to either the HA Appliance's onQ Portal or the DR Appliance's onQ Portal.
2. Go to **DASHBOARD** tab > **RECOVERY NODES** page.

   You see a list of all the recovery nodes and their current status.
3. Locate the recovery node(s).

   Each row in the list represents the recovery node corresponding to a protected node:

Recovery Node		Hostname of the protected node to which this recovery node belongs.
Auto RN Creation	no check mark	Recovery node creation policy is set to On Demand.
	✓ check mark appears	Recovery node creation policy is set to Ready-to-Run.
Space Required	2%	Disk space required to run the recovery node.

Power State	 carat is green and points upward	Recovery node is running and protection is on. Tooltip indicates `RN VM Running`. An agent-less PN is backed up using a PN proxy and the Power State of that PN proxy is detected using the power state of the VM on the ESX Server.
	 carat is red and points downward	Recovery node is not running. Tooltip indicates one of the following: • `RN VM exists but Not Running` – the RN was stopped but not removed and thus the virtual machine exists but is not running. • `RN VM Not Running` – the RN virtual machine was removed or never created.
	 carat is yellow and points upward	Recovery node is starting.
	 carat is yellow and points downward	Running recovery node is shutting down
	 hourglass on a yellow background	Startup dependencies are being evaluated.
Test Mode	no check mark appears	Backup is running in production mode.
	✓ check mark appears	Backup is running in test mode.

Startup Time		When (date and time) the running recovery node was started.
CPU Utilization	field is blank	Recovery node is not running.
	 field is not blank	Recovery node is running. Percentage of the onQ Appliance's allotted CPU used by this recovery node.
Console		Recovery node is running. The icon launches the recovery node's operating system's console.
Appliance CPU Load		Percentage of onQ's CPU usage.
RN Disk Space Available		Total amount of disk space available for RNs. See also Monitor disk space and memory usage.
RN Disk Space Allocated		Total amount of disk space allocated for RNs. See also Monitor disk space and memory usage.
RN Disk Space Used		Total amount of disk space used by RNs. See also Monitor disk space and memory usage.
RN Memory Available		Total amount of memory available to RNs. See also Monitor disk space and memory usage.

Related Topics

Monitor protected nodes

Generate on-demand reports

12.5 Monitor backups

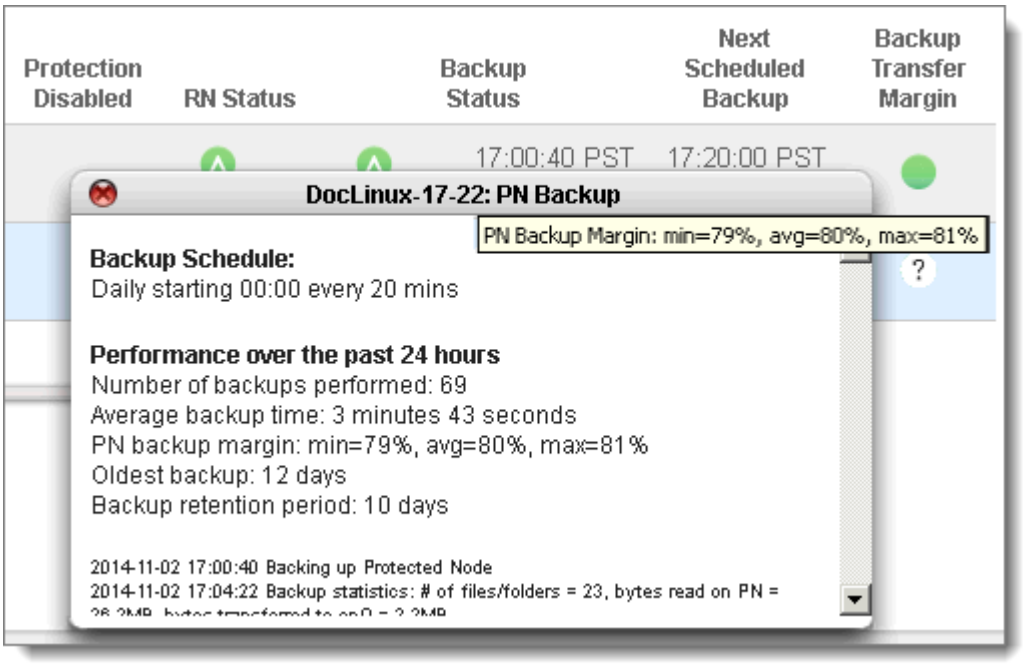
onQ Manager provides two main tools to help you monitor your backups:

- Backup Transfer Margin
- Event Log
- Backup Activity Report
- Alerts

From a protected node's dashboard, you can get insight into backup margin and transfer marge for that protected node's backups. The backup transfer margin icons indicate how well backups and DR transfers are keeping up with the backup interval.

In the event that the interval isn't met, the Event Log indicates `scheduled backup request for <pn_name> occurred before previous backup completed`....In this case, increase your interval time so as to avoid this scenario.

If you want to see recent backup events for individual recovery nodes, simply mouse-over the recovery node on either the HA Appliance or the DR Appliance:



Lastly, be sure to browse the Backup Activity Report on a regular basis and all critical alerts.

Related Topics

Initial Backup
Backup Activity Report
Generate on-demand reports

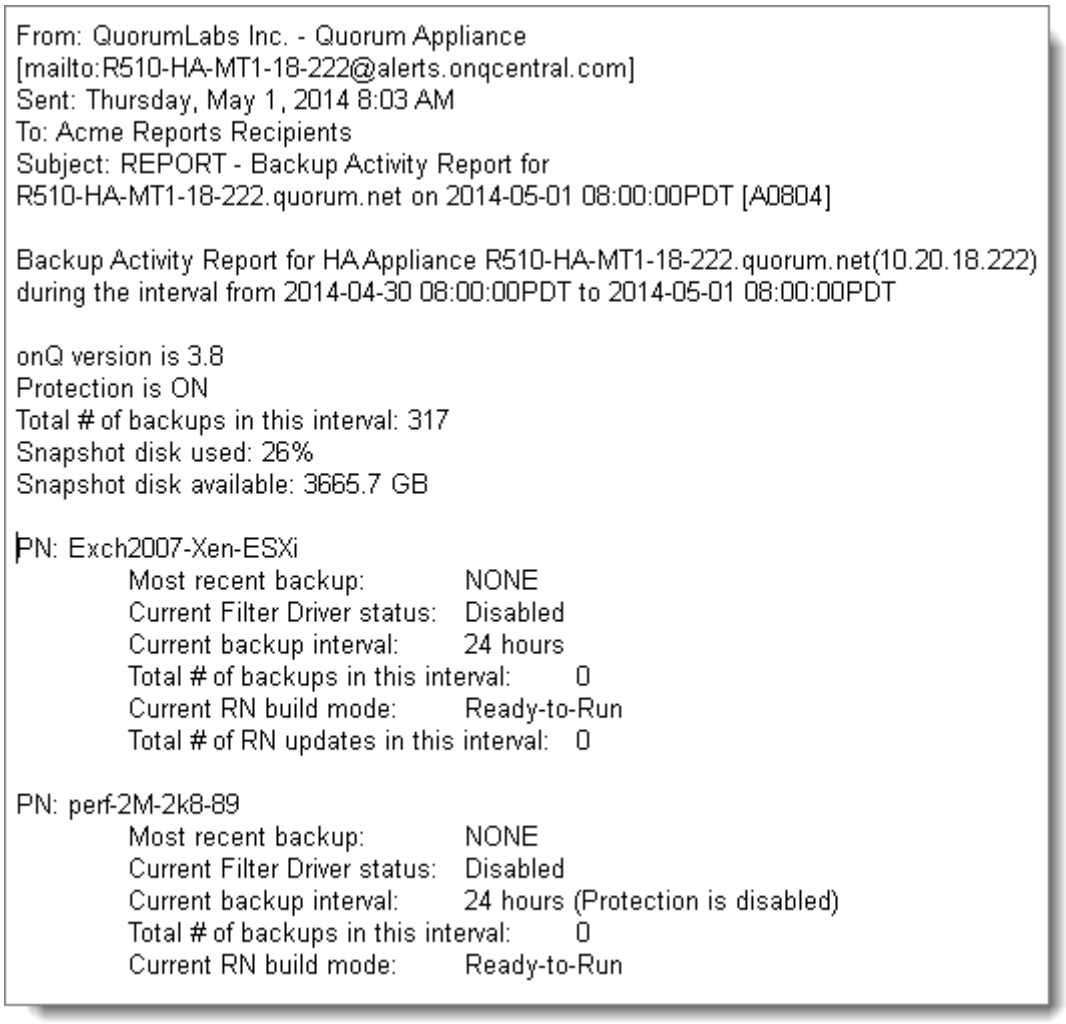
12.6 Backup Activity Report

onQ emails you a Backup Activity Report on a daily basis. If you are not receiving this report, ensure that you have alerts enabled and that your email program is not sending the emails to your SPAM folder.

The Backup Activity Report lists:

- The total number of backups that completed for the day for each protected node and a timestamp for the most recent backup.
- Statistics on the amount of disk space that remains on the Vdisk.

- Status of each PN's Self Test.



From: QuorumLabs Inc. - Quorum Appliance
[mailto:R510-HA-MT1-18-222@alerts.onqcentral.com]
Sent: Thursday, May 1, 2014 8:03 AM
To: Acme Reports Recipients
Subject: REPORT - Backup Activity Report for
R510-HA-MT1-18-222.quorum.net on 2014-05-01 08:00:00PDT [A0804]

Backup Activity Report for HA Appliance R510-HA-MT1-18-222.quorum.net(10.20.18.222)
during the interval from 2014-04-30 08:00:00PDT to 2014-05-01 08:00:00PDT

onQ version is 3.8
Protection is ON
Total # of backups in this interval: 317
Snapshot disk used: 26%
Snapshot disk available: 3665.7 GB

PN: Exch2007-Xen-ESXi
Most recent backup: NONE
Current Filter Driver status: Disabled
Current backup interval: 24 hours
Total # of backups in this interval: 0
Current RN build mode: Ready-to-Run
Total # of RN updates in this interval: 0

PN: perf-2M-2k8-89
Most recent backup: NONE
Current Filter Driver status: Disabled
Current backup interval: 24 hours (Protection is disabled)
Total # of backups in this interval: 0
Current RN build mode: Ready-to-Run

If the Backup Activity Report shows failures, generate the applicable logs to get more information. For example, if a Self Test failed, generate the Self Test Log for the specified time period.

Moreover, the Backup Activity Report has attachments that provide you additional information.

- Backup Statistics Report (`BackupStats.csv`) – byte information for all transfers.
- Alert Activity Report (`Alerts.csv`) – summary of all alerts sent by onQ Central.

- PN Health Report (`Health.txt`) – data to help debug the PN and predict when a given PN will exceed disk space and VSS limits. This report will not appear as an attachment if the report is empty. Also, if your onQ is not configured to send this report to onQ Central, Quorum Support might request it in an effort to debug PN issues. This report is not currently available for agent-less PNs.

Related Topics

Resize protected node's vdisk

About Alerts

Modify e-mail alert settings

12.7 Monitor DR Appliance

Managing disaster recovery with the DR Appliance is slightly different from managing high availability, but just as simple.

The main function of the DR Appliance is to accept PN snapshots from the HA Appliance, vet them, save them to the DR Appliance, and load the most recent snapshot when the DR recovery node is needed.

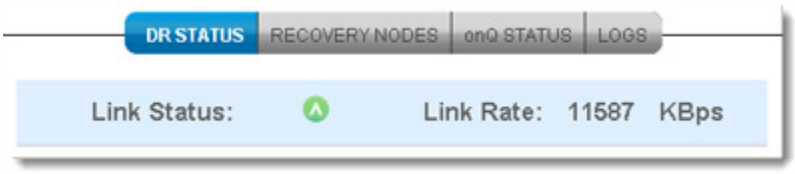
To monitor from the DR Appliance:

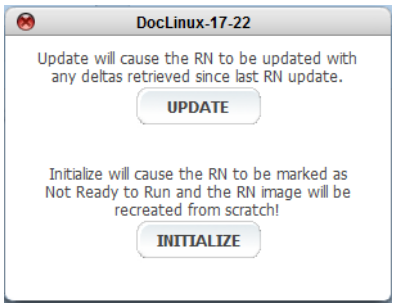
In a nutshell, ensure that:

- the connection between the two onQ Appliance is active.
- each RN's Last Update Time equals the corresponding PN's Backup Status. If not, the RNs and corresponding PNs may be out of sync because (1) The transfer of snapshots from HA to DR paused for some reason, or (2) RNs are not being updated on the DR Appliance.
- there's enough disk space to store the snapshots.

1. Log on to the DR Appliance's onQ Portal.
2. Go to **DASHBOARD** tab > **DR STATUS** page.
3. Observe the following:

Link Status	carat is green and points upward	WAN/VPN connection to the companion HA Appliance is up.
	carat is red and points downward	WAN/VPN connection to the companion HA Appliance is down.
	icon is yellow	Usually indicates that Protection is turned off.



Link Rate		Available bandwidth (in kbps) of WAN/VPN connection to the companion HA Appliance.
RN Name		Hostname of the recovery node. In most cases this hostname is the same as the protected node's hostname. onQ assembles this list of recovery nodes from backup data received from the HA Appliance.
Protection Disabled on HA	no check mark	Protection is enabled on this node.
	✓ check mark appears	Protection is disabled on this node.
RN Ready Status	 carat is green and points upward	Recovery node is ready.
	 carat is green and points upward, and is accompanied by a red asterisk	Recovery node is ready, but it has an `OUT OF DATE RN Template`. You have two choices: 
	 carat is red and points downward	Recovery node is not ready.

RN Running Status	 carat is green and points upward	Recovery node is running and protection is turned on. Tooltip indicates `RN VM Running`.
	 carat is red and points downward	Recovery node is not running and protection is turned off. Tooltip indicates one of the following: • `RN VM exists but Not Running` – the RN was stopped but not removed and thus the virtual machine exists but is not running. • `RN VM Not Running` – the RN virtual machine was removed or never created.
	 hourglass on a yellow background	In production mode, this icon occurs when startup dependencies have been defined for an RN and the RN is being started as part of a group or using the **Start all RNs** button.

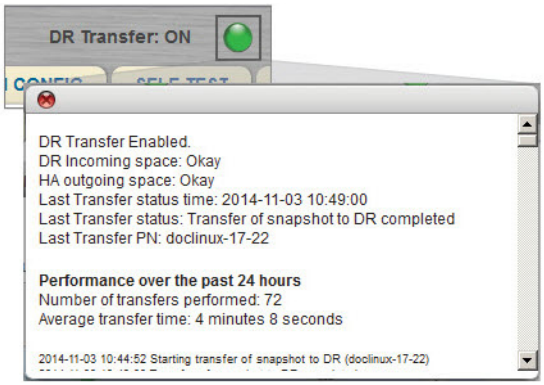
Test Mode	no check mark appears	Backup is running in production mode.
	✓ check mark appears	Backup is running in test mode.
Last Update Time		Last protected node snapshot taken from the DR Repository. RN's Last Update Time must equal the corresponding PN's Backup Status; otherwise the Appliances are not synchronized.
Disk Usage	3 %	Percentage disk usage of entire onQ Appliance, and includes any recovery nodes that are running. To see the files that are consuming this space, go to the DR Repository (click on the **RETRIEVAL** tab).

To monitor from the HA Appliance:

The HA Appliance's onQ Portal has header information that provides information about DR transfers, the HA-DR communication link, and DR disk space.

1. Log on to the HA Appliance's onQ Portal.
2. Go to the **DASHBOARD** tab.
3. Observe the **DR Transfer** status in the page header.

Mouse-over the status icon to view transfer statistics:.

 solid green icon	ON	DR transfers are enabled and the link between the HA and the DR is good. Also, the last transfer, if any, completed successfully. 
 solid red	OFF	DR transfers are disabled.

 solid red	ERR	Either of the following: • the DR returned a link error. • the last transfer to the DR resulted in an error. • protection is off, but starting up. • the disk space on the DR for storing incoming transfers from HA is critically low. • the disk space on HA for storing outgoing transfers to DR is critically low.
 solid yellow icon with question mark	?	onQ is unable to retrieve status information. This status can appear when screen first loads or when protection is off.
 solid yellow with exclamation	WARN	The disk space on the DR for storing incoming transfers from HA is low, or the disk space on HA for storing outgoing transfers to DR is low. This status can also occur if onQ is unable to communicate with the DR.
	XFER, followed by an hour glass ()	A transfer to the DR is underway. This status might not be displayed for all transfers if the transfer is under 30 seconds and the refresh of the status display occurs between transfers.

Related Topics

Monitor backups

Disable replication globally

12.8 Monitor onQ Archive Vault enrollment

After you set up the trust relationship between the onQ Appliance and the onQ Archive Vault (see AV online help for instructions), your onQ Appliance's onQ Portal displays an enrollment status for the AV: the **Enroll** button reads **Disenroll**, and the **Archive Enrollment Status** is `Enrolled`. In addition, the Archive Vault Portal indicates that the onQ Appliance is enabled.

To verify enrollment status:

1. Log on to the onQ Appliance's onQ Portal.
2. Go to **APPLIANCE CONFIG** tab > **ARCHIVE VAULT** page.

   If the onQ is enrolled on the onQ Archive Vault for archiving, then the **Archive Enrollment Status** indicates `Enrolled`.

Related Topics

Enroll and Disenroll onQ Archive Vault

12.9 Monitor disk space and memory usage

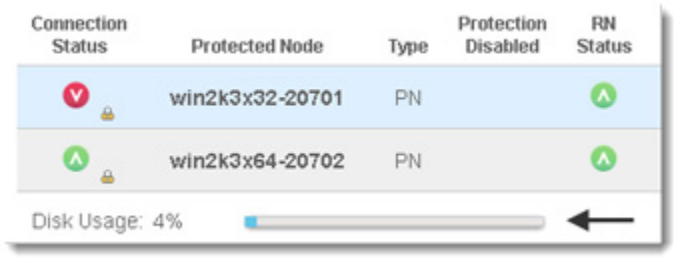
onQ can report on the disk usage of the onQ Appliance's Repository and disk and memory usage of recovery nodes through the onQ Portal as outlined below, or through the onQ Monitor as outlined in Monitor multiple onQ Appliances.

When the onQ Appliance was pre-configured, a specific amount of disk space was allocated to the Repository for snapshots. You cannot change this allocation.

To determine disk space usage of *protected nodes*:

1. Log on to the HA Appliance's onQ Portal.
2. Go to **DASHBOARD** tab > **PROTECTED NODES** page.

A **Disk Usage** bar appears toward the bottom of the page. This bar shows the percentage disk usage on onQ, in two ways: a percentage figure and a corresponding bar graph.

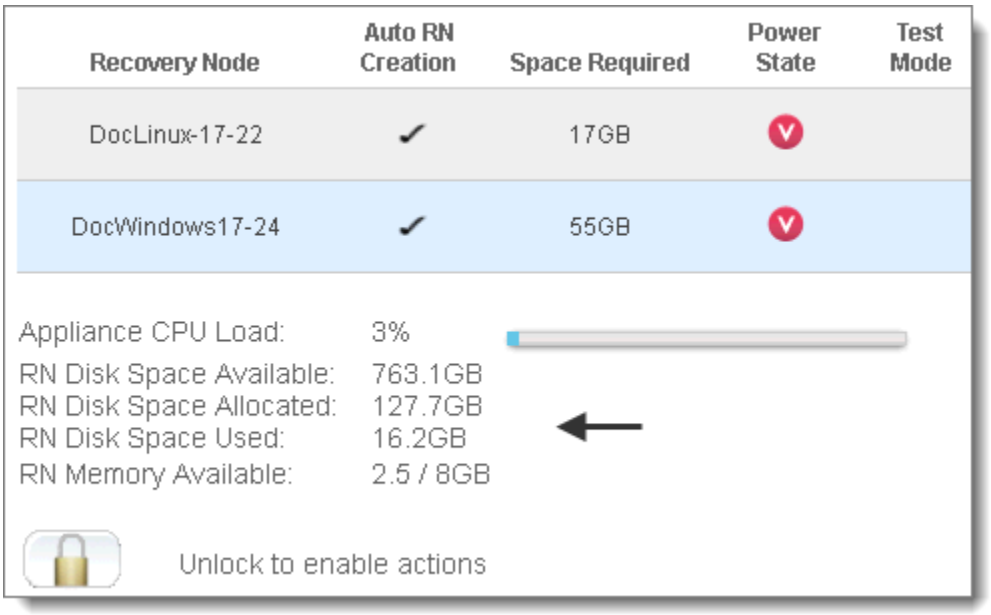


Disk usage is calculated based on the disk usage of all the protected nodes that are running.

To determine disk space usage of *recovery nodes*:

1. Log on to the HA Appliance's onQ Portal.
2. Go to **DASHBOARD** tab > **RECOVERY NODES** page.

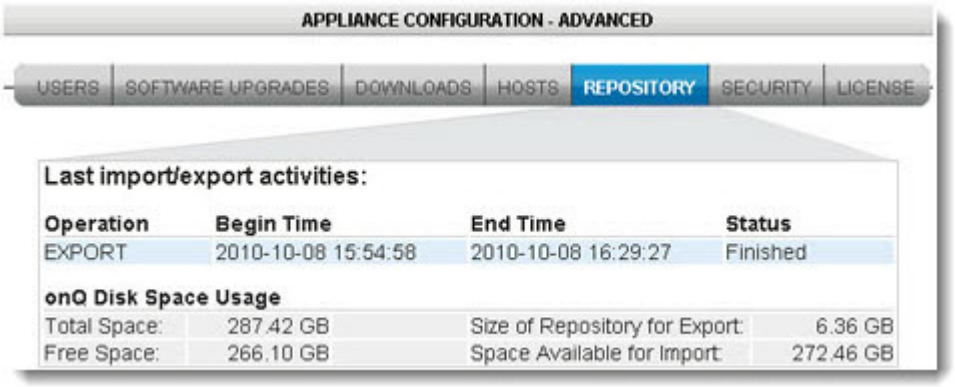
   Statistics appears toward the bottom of the page..



To determine disk space available to the *Repository*:

1. Log on to either the HA Appliance's onQ Portal or the DR Appliance's onQ Portal.
2. Go to **APPLIANCE CONFIG** > **ADVANCED** button > **REPOSITORY** page.

   Free Space indicates the amount of space that the Repository can grow.



Related Topics

Initial Backup

Initialize the DR Repository

12.10 About Alerts

The onQ Manager sends alerts if it detects possible system issues (for example, if a protected node or an incremental snapshot fails). This topic includes the most common alerts that require your attention.

These alerts can be critical or informational only, and Quorum has assigned each alert a preconfigured delivery policy. A summary of alert activity is

available in the Backup Activity Report as an attachment.

Note: If recovery is required, you can start any or all of the recovery nodes. The backup and replication process continues after recovery, based on the recovery nodes and not the original protected nodes.

12.11 Modify e-mail alert settings

onQ sends alerts if it detects possible system failures. To receive such alerts, you must configure the onQ Appliance to send those alerts.

As your alert delivery mechanism, you can specify either onQ Central, which only requires port 443 access to the internet, or e-mail, which requires a SMTP server and, optional credentials.

To modify alert settings:

1. Log on to either the HA Appliance's onQ Portal or the DR Appliance's onQ PortalArchive Vault Portal.
2. Go to **APPLIANCE CONFIG** tab > **ALERTS** page.
3. Click **MODIFY**.
4. In the **Alert Delivery Configuration** pane, do the following:
   a. Select one of the following mechanisms to deliver your alerts.
      - **(Recommended) onQ Central** — Use this option if you want the Appliance to send alerts to onQ Central, Quorum's support management system, in addition to the recipients that you specify. This option provides Quorum Support information about your Appliance as it relates to those alerts, enabling Quorum Support to better respond to potential issues with your Appliance. This option requires that you open outbound port 443 on your firewall.
      - **email** — Use this option if you want the onQ virtual machine to send alerts to use your company's mail server to the recipients that you specify. Specify your company's email server and, optionally, the credentials for authenticating with that server. Ensure that port 25 is open on that mail server.

   b. In the **Daily Report Time** field, specify when you want the report to begin capturing data for the Backup Activity Report. This time is a 24-hr period, and defaults to 8pm.
5. In the **Alert distribution lists** pane, type a comma-separated or space-separated list of email addresses for the recipients:
   - **Reports**. Recipients for the Backup Activity Report.
   - **Notify Alerts**. Recipients for all the notification alerts.
   - **Critical Alerts**. Recipients for all the critical alerts.

6. Click each **Test** button to test your email settings. This test sends a sample email to the recipients that you specified in Step 5.
    - If onQ Central is your delivery mechanism and you don't receive the test email, ensure that port 443 is open on your firewall. In this case, onQ is saving the alerts to its alert queue. onQ will resend these alerts after you configure your firewall correctly.
    - If the onQ virtual machine's delivery mechanism is to an email server (SMTP relay host) such as your company's email server, and you don't receive the test email, ensure that port 25 of that mail server is open to the onQ virtual machine.

12.12 Retrieve logs

Logs provide a convenient record of internal activity (events) on the Appliance. Logs can serve as a useful means of monitoring, compliance reporting, and, if necessary, troubleshooting the onQ Appliance. The log you want to retrieve depends on the activity or component you're trying to monitor or audit.

To retrieve a log:

You can either view or download a given log. You can also use the log filters to view specific data, then export that data to a text file.

1. Log on to either the HA Appliance's onQ Portal or the DR Appliance's onQ Portal.
2. Go to **DASHBOARD** tab > **LOGS** page.

   You see a page similar to the following:
3. From the **Log Type** drop-down list choose one of the following:
    - Event Log — a log of current and past events, listed in alphabetical order by hostname, for all recovery nodes. Use the **Display only errors** check box to generate a log that only includes events that result in failures or errors. This log has no expiration.
    - Event DB Log — a log of database events internal to the Appliance. These events have severity levels. This log is for

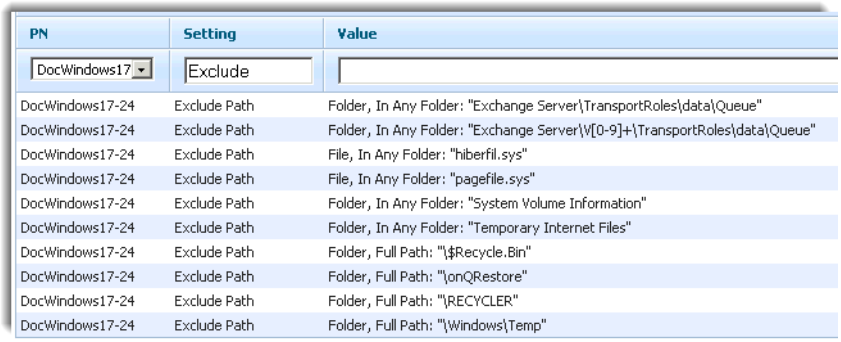
Quorum Support use only.

- Expired Snapshot Log — a log that lists the snapshots and corresponding Protected Nodes that onQ expired and the reason for the expiration. onQ expires snapshots because of a predefined retention policy or critically low disk space.
- FLR Activity Log — a log of file-level restores initiated from the onQ Appliance. The report lists the user that initiated the restore, the time the restore started and completed, the status of the restore, the files and folders that were restored, the source (snapshot), and the machine (PN name or IP address) to which the objects were restored.
- HA --> DR Transfer Log — a log of HA to DR data transfer statistics. The transfer log report includes Protected Node name, start and end times, size of data transferred, and the rate of transfer for each snapshot. The transfer rate is an average, and is calculated based on start/stop times and the size in the log entry.
- Manager Debug Log — a log of debug reports initiated by Quorum. This log is for Quorum Support use only.
- PN Configuration — a report of all the configuration data that is listed in the Protection Config menus, including advanced settings, backup include/exclude lists, startup dependencies, and RN services. This report is often used for compliance reporting. The data is listed by PN, setting, and setting value.
- Self Test Log — a log that shows the results of your RN testing.
- Upgrade Log — a log of all onQ software upgrades for a selected period of time. The report lists the date, time, and the software packages installed on to onQ.
- WSR Activity Log — a log of all Windows Share Restore activity initiated from the Appliance. The report lists the browsing history, file objects accessed, time of access, and the client machine address from which the WSR was initiated.

4. Specify a single PN or all PNs, if applicable.
5. Specify a time range for the data you want to display by specifying the start and end dates, if applicable.
6. Do one of the following:
   - Click **VIEW LOG** to see the information displayed.
   - Click **DOWNLOAD LOG** to download the log to a file on your local computer.

7. (Optional) Use the log filters to include or exclude specific information, then **Save** the filtered data to a text file.

   If you have a browser older than IE 10, this **Save** button doesn't appear. Use the **DOWNLOAD LOG** button instead.



PN	Setting	Value
DocWindows17	Exclude	
DocWindows17-24	Exclude Path	Folder, In Any Folder: "Exchange Server\TransportRoles\data\Queue"
DocWindows17-24	Exclude Path	Folder, In Any Folder: "Exchange Server\V[0-9]+\TransportRoles\data\Queue"
DocWindows17-24	Exclude Path	File, In Any Folder: "hiberfil.sys"
DocWindows17-24	Exclude Path	File, In Any Folder: "pagefile.sys"
DocWindows17-24	Exclude Path	Folder, In Any Folder: "System Volume Information"
DocWindows17-24	Exclude Path	Folder, In Any Folder: "Temporary Internet Files"
DocWindows17-24	Exclude Path	Folder, Full Path: "\$Recycle.Bin"
DocWindows17-24	Exclude Path	Folder, Full Path: "\onQRestore"
DocWindows17-24	Exclude Path	Folder, Full Path: "\RECYCLER"
DocWindows17-24	Exclude Path	Folder, Full Path: "\Windows\Temp"

13

Troubleshooting

- Installation Problems
- (Agent-based PNs) Connection Problems
- (Agent-less PNs) Connection and Backup Problems
- Recovery Node Problems
- (Agent-based PNs) Protected Node Upgrade Problems
- Self-Test and RN Boot Problems
- Helpful System Information
- Generate Support Log
- Backup Alerts
- Snapshot Alerts
- PN Status Alerts
- License Expiration and Upgrade Alerts
- Reboot Alerts
- Protection Alerts
- Recovery Node and PN Disk Space Alerts
- Repository Space
- onQ Disk Space Alerts
- Connection Alerts
- Self-Test Alerts
- Change iDRAC network configuration

13.1 Installation Problems

Use the following table to help you troubleshoot installation problems as reported by the onQ Portal.

Symptom	Possible Cause	Solution
I see a `Browser and onQ are not on the same network` message	This message means that the PN and HA Appliance are on different subnets. This configuration is okay.	If you intend for them to be on different subnets, ignore this message. This message is just a reminder.
(applies to agent-based PNs only) I don't see the **Protect Me** button.	The following conditions must be met before the **Protect Me** button is visible: • The onQ Appliance must have been initially configured. • The PN's operating system must be supported by onQ. • Protection must be turned **OFF**.	Verify that the conditions are met. For os support, go to "Platform Support" in onQ Release Notes.

13.2 (Agent-based PNs) Connection Problems

Use the following table to troubleshoot connection problems with agent-based PNs as reported by the onQ Portal. This table lists the most likely causes first; therefore, verify the possible causes in the order that they are listed.

Symptom: PN's **Connection Status** icon is red	
Possible Cause	**Solution**
System time on the PN is not synchronized with the onQ Appliance. If you're currently installing, this red icon appears 180 seconds into the installation process or 5 minutes after installation. If you've successfully deployed and the PN and the onQ Appliance subsequently become unsynchronized, a missing padlock appears next to the PN, in addition to this red icon.	You have a few choices. Go to Synchronize system time.
PN is down.	Bring up the server.
onQ Service is not running	Restart the service. Go to (Agent-based PNs) Restart the onQ Service.
PN's IP address or hostname changed.	Update IP address and hostname. See Manage hosts.
UDP port 5990 and TCP ports 5000 and 5990 are not open on the PN.	Verify that you performed this step during enrollment: Step 6: Enroll protected nodes.

Symptom: PN's **Connection Status** icon is red	
Possible Cause	**Solution**
Certificate is corrupt.	Regenerate and reinstall the certificate. See Secure connection to PNs.
DNS Server is down	Bring up DNS Server. Add host. See Manage hosts.
Incorrect IP address for DNS Server. If you migrated the DNS Server, onQ might have the incorrect IP address.	Go to the hosts file and update the IP address. See Manage hosts.
If you just performed a restore using QUARK, but the onQ Portal indicates a PN `offline` Connection Status, your PN may have acquired a DHCP IP address, not the original static IP address.	Modify the PN's configuration to reflect its original, static IP address. Afterward, onQ can resume backups of the PN.

Related Topics

Connection Alerts

13.3 (Agent-less PNs) Connection and Backup Problems

Use the following table to troubleshoot connection and backup problems with agent-less PNs as reported by the onQ Portal. This table lists the most likely causes first; therefore, verify the possible causes in the order that they are listed.

Symptom	Possible Cause	Solution
PN's Connection Status icon is red and tooltip indicates `PN offline`.	Firewall might be preventing ping.	Check firewall. Configure it to allow PN to respond to ping.
	PN proxy might be powered off.	On the ESXi host, verify that PN proxy (VM) is powered on. If not, power on the vm.
	If you just performed a restore using QUARK, your PN may have acquired a DHCP IP address, not the original static IP address.	Modify the PN's configuration to reflect its original, static IP address. Afterward, onQ can resume backups of the PN.
PN's Connection Status icon is green with a yellow dash (–) sign and tooltip indicates `Proxy offline`.	PN proxy might be in the process of restarting (every backup request from onQ does so).	Wait for the PN proxy to restart.
	PN proxy might be in a bad state.	Try re-installing via portal enrollment.
	Unknown.	Log on to the ESXi host and manually reset the PN proxy's power.

Symptom	Possible Cause	Solution
Backup process fails to create snapshots and the PN's Event Log indicates an event for this issue.		Disable backups for the PNs from the same host, then initiate a manual snapshot. If the snapshot completes successfully, delete that snapshot.
		From the ESXi host, reset power on the PN proxy.
		Restart protection on the onQ.
Backup of PN's pauses and the PN's Event Logg indicates that there is a backlog for backup requests.		From the ESXi host, power off the PN proxy (VM). onQ restarts the proxy and starts backups again.

Symptom	Possible Cause	Solution
PN proxy returns the following exception: `system: 2014-06-02 06:31:29: Snapshot of PN and mount to proxy failed: PN wasn't moved to the new snapshot, or Snapshot creation failed.`		Base-disk must be mounted within 30 minutes. VMware supports creation of snapshots via the vSphere SDK. The PN proxy uses this API to create a snapshot and mount the snapshot disks. In some cases, (for example, a faulty VSS in the PN or a quiesce script in the PN), the snapshot API can fail. After the vSphere SDK returns a success for snapshot creation, the proxy waits 30 minutes for the base-disk to be mounted. If the snapshot disk isn't mounted during that time and the PN hasn't moved to the new snapshot disks, the PN proxy fails the operation and errors out with aforementioned exception.

13.4 Recovery Node Problems

Use the following table to troubleshoot problems with your recovery nodes as reported by the onQ Portal.

Symptom	Possible Cause	Solution
RN's RN Status icon is red	You just added the PN and onQ is backing up the PN for the first time.	Wait for the backup to complete. See also Monitor backups.
	Protection is off.	Turn on protection. Go to Start node protection.
	Creation policy is set to Build on Demand.	Change the policy. Go to Modify recovery node creation policy.
Red exclamation point appears next to the protected node	The protected PN is down and the corresponding RN is running.	This is expected behavior. Fix your PN. Before you bring the PN online, stop the RN.

Related Topics

Recovery Node and PN Disk Space Alerts

13.5 Self-Test and RN Boot Problems

Use the following table to troubleshoot self-test failures with agent-based and agent-less PNs as reported by the onQ Portal and one or more of the alerts outlined in Self-Test Alerts. For each symptom, this table lists the most likely cause first; therefore, verify the possible causes in the order that they are listed.

Before You Begin: Verify that the RN is in a usable state by booting it in test mode. If the RN fails to boot, reinitialize it to eliminate the possibility that the problem is with the RN build process itself. Afterward, proceed with the table

below.

Symptom	Possible Cause	Solution
RN may or may not boot	Occurs immediately following a Linux PN enrollment. Typo in boot menu loader or mismatch between the kernel versions and the initramfs version.	If you're in the process of enrolling the Linux PN for the first time and the RN won't boot, see the troubleshooting tips in (Agent-based Linux PNs) Enroll protected nodes or (Agent-less Linux/Windows PNs) Enroll protected nodes.
	Operating system is not supported.	Verify that the operating system is supported. Go to "Platform Support" in onQ Release Notes.
	RN has insufficient resources.	If RN boots, verify that there is adequate memory and disk space for the RN: go to Monitor disk space and memory usage. If there isn't, provide the RN more resources.
	Unresolved dependencies led to the onQ Service not starting on time on the RN.	Determine if the RN is dependent on another service, then make that service available. Some services might take time to start; eventually those services might succeed or fail or time out thereby delaying the onQ Service from starting on time. Such services might rely on other resources (internal or external). For example, the service might be waiting for a mount manager (internal resource) to process huge mount points leftover in the registry, or for a time server/Domain Controller (external resource) to become available.

Symptom	Possible Cause	Solution
RN booted with correct IPs but no network activity	PN has a dynamic IP address.	Make sure that the PN has a static IP.
	PN is running on an OEM version	Make sure that the PN is not running an OEM version. Networking might be disabled. For more information, go to Go to "Platform Support" in onQ Release Notes.
BSOD and RN cannot boot	RN has a faulty service	If there is a blue screen, read the blue screen code as it might indicate the service that caused the blue screen. Disable that faulty service; if that doesn't work, capture the BSOD screen and contact Quorum Support.
RN/PN takes too long to boot	RN has slow-to-start service	Does the RN take more than 15 minutes to come up? If yes, check the Windows event log and figure out which service(s) is taking a long time to start or which service(s) is timing out. Disable or reconfigure this service to speed up the boot time.
	PN has large registry	Check the PN's boot time. If the PN has a long boot time, the problem needs to be addressed at your PN site. For example, the PN might have a large registry that causes the PN to boot slowly. Resolve the root cause of the growing registry and clean up the registry.

Symptom	Possible Cause	Solution
RN reboots continuously	Windows Server 2012 R2 Bug	(Windows Server 2012 R2) If RN continues to reboot, boot it in self-mode, then search the Windows event logs for the suspicious service and disable it. For example, the `ShellHWDetection` service is notorious for continuously rebooting an RN running Windows Server 2012 R2. Disabling the `ShellHWDetection` service resolves the problem.
RN boots, but has no XenServer NIC hardware	3rd party software	In specific cases, some 3rd party software applications prevent the XenServer NIC from functioning. For example, Symantec endpoint protection 12.1.4100.426 blocks the XenServer NIC from functioning. In this case, disabling the Symantec endpoint protection services resolves this problem.

Symptom	Possible Cause	Solution
RN boots with XenServer NIC, but has an incorrect IP address	Misconfigured RN network	Check the RN NIC configuration and IP configuration via the onQ Portal. Verify that `xvf.dat` exists on the PN and that the PN does not have a custom network configured. Make sure the network configuration is correct.
RN boots with XenServer NIC, but has no IP address	Misconfigured cluster network	RN does not own the cluster resources as set forth in the cluster policy, so the IP address is inactive. If you want the RN to own the cluster resources, use the Enable Cluster Support parameter or change the cluster policy.
	3rd party software	In specific cases, some 3rd party software applications prevent the IPs from attaching to XenServer NIC. For example, Network load balancing software (NLB) might block an IP from being used. Also, the WLBS Windows NLB service blocks NIC IPs on Windows Server 2012R2; therefore, delete the WLBS service on the RN build via the onQ Portal, if required. Disable the WLBS service on other Windows Server platforms as needed.
RN booted with correct IPs but no network activity or limited network connectivity	3rd party software	Check if any 3rd party network management software is installed: `Antivirus/firewall/network manager/...`. If yes, disable this software on the RN build.

If the self-test still fails or the RN cannot boot, contact Quorum Support, providing the information outlined in Helpful System Information. In addition, provide us the answers to the following questions to speed up resolution:

- Is there a PN with the same operating system on the same onQ Appliance where self-tests are passing?
- Is this a newly protected PN?
- Is this PN experiencing *sudden* self-test failures with prior success?

(Important!) In addition:

- Did you provide a screen shot of RN's Programs and Features (Control Panel > Programs and Features)? If not, please do so. This output tells Quorum what 3rd party software is installed.
- Did you include a screen shot of the RN's BSOD, if applicable? If not, please do. This output might indicate the offending service.
- Did you include RN's output from `systeminfo` command? If not, please do.

13.6 (Agent-based PNs) Protected Node Upgrade Problems

Use the following table to troubleshoot problems with your protected nodes during or following an upgrade.

Symptom	Possible Cause	Solution
You see the following error during a PN upgrade following an onQ upgrade and protection restart: `Error updating service on many 2003 PNs (msi: code 800b0101)`	The global root certificate on your Windows PN is invalid or expired. The upgrade is not the cause; moreover, any industry software upgrade requires a valid global root certificate.	Do the following: 1. Install a new global root certificate on the PN from https://support.globalsign.com/customer/portal/articles/1426602-globalsign-root-certificates. 2. Perform a Windows update.

Related Topics

PN Status Alerts

13.7 Helpful System Information

To speed up technical support for your issue, please provide the following information when contacting Quorum Support.

- Account information.
- PN's operating system, host name, hardware vendor/model, physical or virtual: # `systeminfo`.
- onQ STATUS page—screen shot or copy/paste.

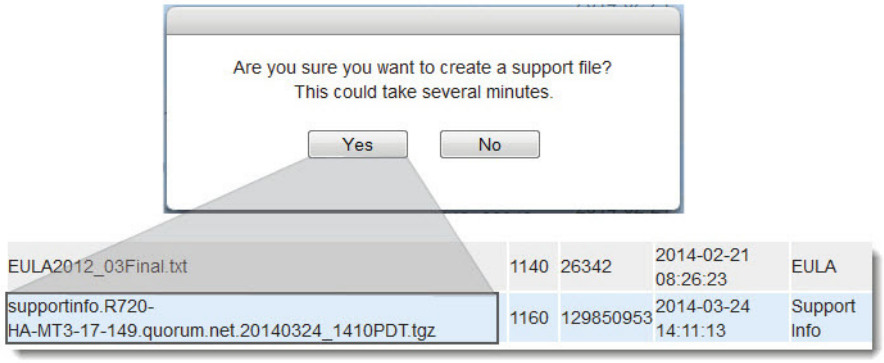
- Screen shot of special networking configuration and software installation: Control Panel > Programs & Features.
- BSOD screen shot, if applicable.
- List of any recent upgrades or software changes.
- If your PN is VMware-hosted, specify whether it's agent-based or agent-less.
- Support Log (see Generate Support Log).
- The Health Report, if your onQ is not configured to send these reports to onQ Central.
- List of your troubleshooting steps and results.

13.8 Generate Support Log

If requested by Quorum Support, email your support engineer the Support Log. This log provides additional information to help Quorum Support troubleshoot issues with your onQ Appliance.

To generate a Support Log:

1. Log on to the onQ Appliance's onQ Portal.
2. Go to **APPLIANCE CONFIG** > **ADVANCED** > **DOWNLOADS**.
3. Click on the **SUPPORT** button, then **Yes** to generate the log (`.tgz` file).
4. From the **DOWNLOADS** tab, scroll down to the end of the download list and select the file that you just generated, then click the **DOWNLOAD** button to save the Support Log.



5. Email this Support Log to Quorum Support.

   This Support Log is quite large, and many email systems will reject it. If you have problems emailing this file, contact Quorum Support.

13.9 Backup Alerts

It's important that you resolve backup failures. Fortunately, most failures can be attributed to are few common causes, including conflicting VSS-based (Volume Shadow Copy Service) backup solutions, insufficient disk space, and high CPU usage.

For specific alerts:

- A0801
- A0802
- A0803

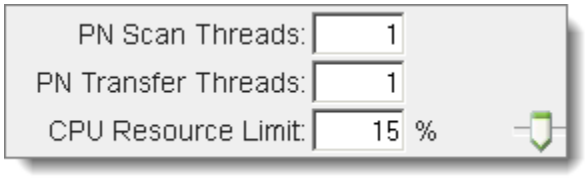
13.9.1 A0801

Message: The onQ backup of Protected Node *PNName* has repeatedly failed. The last backup attempt returned code: *ReturnCodeListedBelow*. Please verify that there is sufficient free disk space on the Protected Node. Backups will continue on the configured schedule or when triggered from the onQ web portal.

Cause: The specific cause depends on the return code that appears in the alert message.

Solution: Regardless of your return code, perform the following solutions in

the following table and in the order listed:

Possible Cause	Solution
High CPU Usage	Go to the Task Manager and determine if you have high CPU usage (> 30%). If yes, reduce the load on the system. In addition, go to **PROTECTION CONFIG** tab, select the PN, **MODIFY** > **ADVANCED**, then do the following: **1.** Reduce scan threads to 1. **2.** Reduce transfer threads to 1. **3.** Reduce CPU resource limit by 50% of current limit. The system default is 30%. Set it to 15%. 
Conflicting VSS-based Backup Solutions	There are conflicting VSS-based backup solutions running. Disable all conflicting VSS-based backup solutions.
Insufficient Disk Space	Backups can fail if the PN is running out of disk space. Free up some space on the PN by deleting temp files or, if the PN is a virtual machine, expand the disk.
Other Solutions	Restart the onQ Service and Restart protection.

If these solutions don't work, perform the solution(s) outlined in the applicable return status listed under the drop-down hotspots below and in alphabetical order.

\\?\GLOBALROOT\Device\HarddiskVolumeShadowCopy3\ (The system cannot find the path specified.

There are conflicting VSS-based backup solutions running: see Conflicting VSS-based Backup Solutions.

Block sums file c:\Program Files\Quorum\usr/blocksums/99restart-

blocksums.bin corrupted: expected file object

Delete all files under blocksums folder.

DoSnapshotSet

Do the following:

1. From the PN, run the following command: .

```
# vssadmin list writers
```

2. If writers are at a failed state, reboot the PN.

If that solution doesn't work, contact Quorum Support.

Error writing raw data to c:\Program Files\Quorum\usr/block-sums/sorted-blocksums.bin (Broken pipe)

PN is running out of disk space: see Insufficient Disk Space. Also, delete all files under blocksums folder.

Error writing c:\Program Files\Quorum\usr/scan/0.txt (No space left on device)

The PN ran out of space during a backup. Free up some space: see Insufficient Disk Space.

Error: E_INVALIDARG

Reboot the PN. If that solution doesn't work, reinstall the node software as outlined in (Alternative) To manually install or reinstall agent-based Windows PN software:

Error: VSS_E_VOLUME_NOT_SUPPORTED (..\volumes_win32.cpp @ 297)

We do not support non-NTFS volumes as outlined in Limitations.

Fail to rotate log, QFilter

Do the following:

1. From a cmd prompt, run the following commands:

```
# net stop quorum
# sc stop qfilter
```

2. From `C:\Program Files\Quorum\usr`, delete all files whose filenames start with `qfilter.log`,
3. Run the following command:

```
# net start quorum
```

Invalid Drive

The drive needed to perform the backup is no longer available. Log on to the PN and verify (Windows Explorer > My Computer) that all the drives are available. If the drives are present, reboot the PN. If the drives are missing, add the drives.

NotStarted

There are conflicting VSS-based backup solutions running: see Conflicting VSS-based Backup Solutions. If that solution doesn't work, reboot the PN.

Oracle VSS Writer failed during snapshot, aborting

Restore the Oracle database using one of the following procedures:

- Back up and restore Oracle 11g database on Linux
- Back up and restore Oracle 10g+ database on Windows

A Windows PN backup fails if onQ detects a VSS_WS_FAILED_AT_*ErrorState* for any of the following VSS Writer error states, and records the error state in the PN's Event Log.

`VSS_WS_FAILED_AT_IDENTIFY`

`VSS_WS_FAILED_AT_PREPARE_BACKUP`

`VSS_WS_FAILED_AT_PREPARE_SNAPSHOT`

`VSS_WS_FAILED_AT_FREEZE`

`VSS_WS_FAILED_AT_THAW`

`VSS_WS_FAILED_AT_POST_SNAPSHOT`

`VSS_WS_FAILED_AT_BACKUP_COMPLETE`

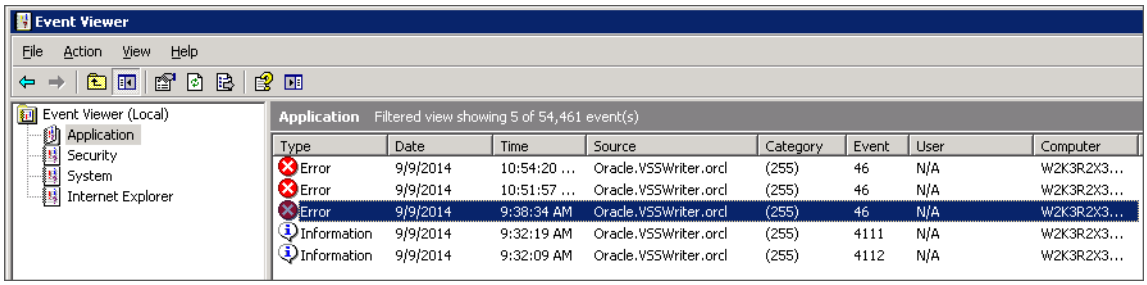
`VSS_WS_FAILED_AT_PRE_RESTORE`

`VSS_WS_FAILED_AT_POST_RESTORE`

`VSS_WS_FAILED_AT_BACKUPSHUTDOWN`

To troubleshoot these error states, look at the PN's Windows log via Event

Viewer: **Start** button > **Control Panel** > **System and Security** > **Administrative Tools** > **Event Viewer**.



For an explanation of the error states, go to Microsoft's Volume Shadow Copy API Reference documentation. For overview information, go to Microsoft's Event and Error Handling Under VSS. Contact your Oracle database administrator to resolve the error.

scan failed, aborting

The PN might not have enough disk space: see Insufficient Disk Space.

If that solution doesn't work, do the following:

1. From the PN, run the following command: .

```
# vssadmin list writers
```

2. If writers are at a failed state, reboot the PN.

If that solution doesn't work, there might be conflicting VSS-based backup solutions running: see Conflicting VSS-based Backup Solutions.

If that solution doesn't work, reboot the PN.

SS_E_PROVIDER_VETO

There might be too much CPU load: see High CPU Usage. If that solution doesn't work, there might be conflicting VSS-based backup solutions running, resulting in VSS writer failures: see Conflicting VSS-based Backup Solutions. If that solution doesn't work, reboot the PN.

Transfer incomplete

If your PN is a Windows Server 2003, you might need to change the values for the `PoolUsageMaximum` and `PagedPoolSize` registry settings as

outlined in Microsoft KB 304101.

Unable to resolve target hostname

Do the following:

- Verify that the hostname hasn't changed. If it has, update the IP address in the onQ's hosts file as outlined in Manage hosts.
- Verify that the DNS entry is correct. If it isn't, update the entry in the onQ's hosts file as outlined in Manage hosts.
- If this PN is in a cluster, ensure that you configured things correctly. Go to Enrollment in Windows Cluster Services Environment.

vss unexpected provider

Oracle VSS Writer failed during a snapshot. There are conflicting vss backup solutions running. *Uninstall* any conflicting VSS-based backup solutions. If this doesn't work, reboot the PN.

VSS_E_SNAPSHOT_SET_IN_PROGRESS

There might be conflicting VSS-based backup solutions running: see Conflicting VSS-based Backup Solutions. If that solution doesn't work, then reboot the PN.

13.9.2 A0802

Message: The onQ backup of Protected Node *PNName* has failed due to a failure of the PN. Backups will continue when the PN comes online.

Cause: A few issues can cause this problem: (1) PN and onQ do not have the same system time; (2) High CPU load; and (3) Other connectivity problems between the PN and onQ.

Solution: Verify system time on both the PN and onQ. If necessary, update the system time. Go to Synchronize system time.

If that solution doesn't work, verify connectivity by attempting to ping the onQ from the PN (see also (Agent-based PNs) Connection Problems).

If that solution doesn't work, you might have high CPU load: see High CPU Usage.

13.9.3 A0803

Message: The onQ service on *PNName* is unable to stop the backup process, `squirtcopy.exe`. This may be due to a transient problem with `taskkill.exe`. Backups on *PNName* have been suspended until protection is stopped and restarted.

Cause: The cause depends on a number of factors.

Solution: Restart protection. If the problem persists, restart the PN.

13.10 Snapshot Alerts

Snapshot issues are rare. However, when they occur, a rescan or protection restart typically resolves the problem.

For specific alerts:

- A0401
- A0402
- A0403
- A0404

13.10.1 A0401

Message: *ApplianceName* has detected an unexpected error during the BCVPreAdd process for *PNName*. The return status was: *ReturnStatus*. Protection will be disabled for this PN.

Cause: The cause depends on a number of factors.

Solution: Restart protection. If that solution doesn’t work, contact Quorum Support.

13.10.2 A0402

Message: *ApplianceName* has detected that an inconsistent object count has been stored in the repository. The HA onQ should have sent an alert for 'Object count inconsistency'. Please refer to that alert for information on resolving this issue.

Cause: The cause depends on a number of factors.

Solution: From the HA Appliance (always start with the HA, even when resolving DR error messages):

1. Force a full rescan. Go to Perform full rescan on PN.
2. Trigger an immediate backup. Go to Initiate immediate backups.

The database will resync. After the full snapshot is processed on the HA, the error on the HA will go away. After the same snapshot is transferred and fully processed/consumed by the DR, the error message on the DR will also go away. Because of the time needed to transfer and process the snapshot, there will be a delay between when the HA Appliance resolves the issue and when the DR Appliance resolves the same issue.

13.10.3 A0403

Message: *ApplianceName* has detected that the object count for *PNName* is inconsistent with the value stored in the repository. A full rescan of the PN is required to resolve this issue.

Cause: The cause depends on a number of factors.

Solution: From the HA Appliance (always start with the HA, even when resolving DR error messages):

1. Force a full rescan. Go to Perform full rescan on PN.
2. Trigger an immediate backup. Go to Initiate immediate backups.

The database will resync. After the full snapshot is processed on the HA, the error on the HA will go away. After the same snapshot is transferred and fully processed/consumed by the DR, the error message on the DR will also go away. Because of the time needed to transfer and process the snapshot, there will be a delay between when the HA Appliance resolves the issue and when the DR Appliance resolves the same issue.

13.10.4 A0404

Message: *ApplianceName* has detected an unexpected error during the BCVAdd process for *PNName*. The error status was: *ReturnStatus*. Protection will be disabled for this PN until protection is restarted for this onQ appliance.

Cause: The cause depends on a number of factors.

Solution: Restart protection. If that solution doesn't work, contact Quorum

Support.

13.11 PN Status Alerts

If the PN (or HA Appliance) is actually down, perform the appropriate DR workflow in conjunction with your business continuity plan:

- (Workflow) Fail over HA Appliance to DR Appliance
- (Workflow) Fail back DR to HA Appliance
- (Workflow) Fail over a PN to an RN
- (Workflow) Fail back an RN to a PN

If the PN is up, the problem can usually be resolved by updating the system time or the onQ hosts file.

For specific alerts:

- A0704
- A0710

13.11.1 A0704

Message: *ApplianceName* has detected a FAILURE of the Protected Node *PNHostName*. If this is an actual failure, use the onQ web portal to start the backup Recovery Node for this machine.

Cause: Either a host IP changed or, more commonly, the system time being off by roughly 5 minutes. Typically when a PN shows red for offline it is one of these 2 cases. However, this error can also occur if the PN shut down or rebooted or the onQ Service isn't running.

Solution: Do one of the following, and in the following order:

- *Time off Scenario*: Verify system time on both the PN and onQ. If necessary, update the system time. Go to Synchronize system time. Finally, restart protection. The PN should reconnect to the onQ.
- *IP Address Scenario*: This problem most commonly occurs when the onQ Service is installed during a PN enrollment when the PN has a DHCP address, prior to a static being assigned. If the IP address of the PN

changed, update the IP address in the onQ's hosts file as outlined in Manage hosts, then restart protection.

- *onQ Service not running on agent-based PN*. PN is up and running, but the onQ Service isn't. This situation is possible following a successful automatic upgrade of the PN or after a reboot of the PN. If true, restart the onQ Service.
- *PN shutdown or reboot*. The PN shut down or rebooted for maintenance or some other reason. If true, power on the PN or wait for the PN to reboot, then verify that the PN is online.
- *Actual PN failure*: If true, go to (Workflow) Fail over a PN to an RN.

13.11.2 A0710

Message: *ApplianceName* was unable to upgrade the onQ Service on *PNName* to version *VersionNumber*. The current version on the PN is *VersionNumber*. The return status from the upgrade attempt was: *ReturnStatus.*

Cause: The most likely cause for this agent-based PN upgrade is that the PN has too many other programs running thereby causing high CPU utilization, or the installer is currently installing a different program.

Solution: Fix the utilization problem. If that solution doesn't work, reinstall the node software as outlined in (Alternative) To manually install or reinstall agent-based Windows PN software: or (Agent-based Linux PNs) Enroll protected nodes.

13.12 License Expiration and Upgrade Alerts

Each Appliance requires a license. Without a valid license, your nodes are not protected. onQ sense license alerts before and after expiration.

For specific alerts:

- A0008
- A0009
- A0014
- A0015

- A0017

13.12.1 A0008

Message: The installed license has expired. Protection on this onQ appliance has been disabled and will not run until a new license is installed.

Cause: The license expired.

Solution: Upload a replacement license. Go to Install onQ Appliance license.

13.12.2 A0009

Message: The installed license has expired. Protection on this Archive Vault appliance has been disabled and will not run until a new license is installed. Please send an email to license@quorum.net with your appliance serial number if you have already purchased your new license.

Cause: The license expired.

Solution: Upload a replacement license. Go to Install onQ Appliance license.

13.12.3 A0014

Message: A software version mismatch has been detected between this appliance and *ApplianceName*. The current running version for this appliance is *VersionNumber*, while the version reported by *ApplianceName* is *VersionNumber*. Due to compatibility reasons no further configuration changes, such as newly enrolled protected nodes, will be available on this appliance until this issue has been resolved. When possible please ensure each appliance is updated to the same version level.

Cause: Each of the onQ Appliances listed are running different software versions. They must be in sync.

Solution: Compare the versions between the onQ Appliance, then resolve the incompatibility by upgrading each onQ Appliance to the same version level. If the documented upgrade process does not support your upgrade path, contact Quorum Support.

13.12.4 A0015

Message: A software version mismatch has been detected between this appliance and *ApplianceName*. The current running version for this appliance is *VersionNumber*, while the version reported by *ApplianceName* is *VersionNumber*. Due to compatibility reasons no further configuration changes, such as newly enrolled protected nodes, will be available on the remote appliance until this issue has been resolved. When possible please ensure each appliance is updated to the same version level.

Cause: Each of the onQ Appliances listed are running different software versions. They must be in sync.

Solution: Compare the versions between the onQ Appliance, then resolve the incompatibility by upgrading each onQ Appliance to the same version level. If the documented upgrade process does not support your upgrade path, contact Quorum Support.

13.12.5 A0017

Message: The installed license will expire on *Date*. You must install a new license before that date in order to maintain protection.

Cause: The license is about to expire.

Solution: Upload a replacement license. Go to Install onQ Appliance license.

13.13 Reboot Alerts

If onQ identifies a problem, determines that a reboot will fix the problem, onQ initiates that reboot and restarts protection. If onQ has enough time prior to the reboot to log the cause for the reboot, it will do so, providing a detail Cause in the alert. Otherwise, a generic message appears in the alert.

For specific alerts:

- A0010
- A0011
- A0012

• A0013
• A0018

13.13.1 A0010

Message: Protection is enabled but the manager process cannot be found. Rebooting the onQ VM and restarting protection.

Cause: This is expected behavior.

Solution: Nothing. onQ rebooted the onQ VM and restarted protection. However, If protection does not start correctly, contact Quorum Support.

13.13.2 A0011

Message: Protection is enabled but the manager process cannot be found. Rebooting the Archive Vault VM and restarting protection.

Cause: This is expected behavior.

Solution: Nothing. onQ rebooted the onQ VM and restarted protection. However, If protection does not start correctly, contact Quorum Support.

13.13.3 A0012

Message: Rebooting the onQ VM and restarting protection to clear 3761.

Cause: This is expected behavior.

Solution: Nothing. onQ rebooted the onQ VM and restarted protection. However, If protection does not start correctly, contact Quorum Support.

13.13.4 A0013

Message: Rebooting the Archive Vault VM and restarting protection to clear 3761.

Cause: The cause depends on a number of factors.

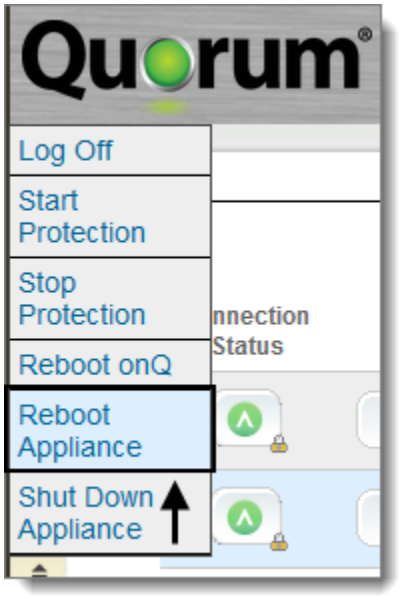
Solution: Nothing. onQ rebooted the onQ VM and restarted protection. However, If protection does not start correctly, contact Quorum Support.

13.13.5 A0018

Message: Unable to shut down *VMName* VM while starting onQ protection.

Cause: The cause depends on a number of factors.

Solution: If an RN is in production due to a disaster, do not reboot the onQ Appliance; contact Quorum Support instead. If no RN is in production, go ahead and reboot the onQ Appliance.



13.14 Protection Alerts

The following alerts are sent out when protection is turned on or off on the onQ by a user that has VARAdmin user role.

- A0001
- A0003

13.14.1 A0001

Message: onQ protection has been stopped at user request

Cause: A user that has VARAdmin user or Admin user role stopped

protection on the protected node.

Solution: Nothing as there are cases in which you might want to stop protection. However, don't forget to start protection afterward; otherwise, your nodes will not be protected. Also, if you did not perform this action, identify other administrators who have credentials to access the onQ.

13.14.2 A0003

Message: onQ protection has been started at user request

Cause: A user that has `VARAdmin` user or Admin user role started protection on the protected node.

Solution: Nothing as you must start protection in order to protect your nodes. However, if you did not perform this action, determine who is managing the onQ.

13.15 Recovery Node and PN Disk Space Alerts

A protected node's virtual disk (vdisk) requires a specific amount of disk space. onQ alerts you when your protected node's vdisk capacity usage exceeds 90%.

For specific alerts:

- A0901
- A0902
- A0903
- A0904
- A0906

13.15.1 A0901

Message: An unexpected error was encountered when attempting to update the contents of a Recovery Node. The error status returned was: *ReturnStatusListedBelow*. No further updates to the Recovery Node will be attempted until you select the Initialize button.

Solution: Depends on the return status listed under the drop-down hotspots

below.

InvalidData

Do not reinitialize the RN. This issue might indicate a problem with the Repository. Contact Quorum Support.

SelectTimeout

Reinitialize the RN.

SystemErrDiskFull

Data in one or more protected volumes has exceeded its virtual disk size allocated on the onQ. Resize the vdisk for the volume(s) indicated in the alert.

13.15.2 A0902

Message: Volume *DriveLetter* of Recovery Node *RNName* is configured with disk space insufficient to create a new RN from the current snapshot. The previously built RN left unchanged. The snapshot data for the PN is intact and a new RN will be created after the configured disk space is increased. You should increase the amount of disk space allocated to this RN.

Cause: Initial vdisk size allocated for a volume on the PN is inadequate because of the growing data on the PN in the corresponding volume. In this case, the onQ Manager does not update the RN.

Solution: Resize the vdisk 15–20% over what is currently being used.

13.15.3 A0903

Message: The Recovery Node disk space allocated for volume(s) *DriveLetter* of *RNName* is over 90% used. The RN will be created, however you may wish to increase the amount of disk space allocated to each named above volume of that RN.

Cause: Initial vdisk size allocated for a volume on the PN is inadequate because of the growing data on PN in the corresponding volume. In this case, the onQ Manager does not update the RN.

Solution: Resize the vdisk 15–20% over what is currently being used.

13.15.4 A0904

Message: The R2V Export process on *ApplianceName* has returned an unexpected error code. This is preventing the update of the RN for *RNName*. The error code return was: `ExportFailed`.

Cause: This is a transient error.

Solution: Wait for one more update to occur. Contact Quorum Support if next update is not successful.

13.15.5 A0906

Message: The onQ restriction [no more than 15 virtual disks per Recovery Node] was hit while making a copy of volume $arg1: Reinitializing the RN image may reduce the number of virtual disks involved.

Cause: You exceeded a virtualization limitation (see Resize protected node's vdisk for explanation).

Solution: Rebuild the RN. If a rebuild doesn't work, contact Quorum Support.

13.16 Repository Space

The Repository needs a certain amount of disk space. If you're low on disk space, there are things that you can do to free up space before the issue becomes critical.

For specific alerts:

- A0101
- A0102
- A0103
- A0104
- A0303
- A0304

13.16.1 A0101

Message: The free Repository space on *ApplianceName* has fallen below normal limits. If the free Repository space continues to decrease, onQ will expire the oldest backups for each Protected Node in order to recover space.

Cause: This alert is for informational purposes: repository disk utilization is at 75%. onQ has a built-in policy whereby if its repository disk space utilization exceeds 85% onQ begins to "free up" space by deleting PN backups, starting with the oldest backups.

Solution: Delete orphan backups manually, reduce retention period, or remove a node for which you don't need protection. Contact Sales to discuss options for expanding disk space.

13.16.2 A0102

Message: The free Repository space on *ApplianceName* has fallen below normal limits. If the free Repository space continues to decrease, the Archive Vault will disable the addition of any new archive jobs until additional space becomes available.

Cause: This alert is for informational purposes: repository disk utilization is at 75%. AV has a built-in policy whereby if its repository utilization exceeds 85%, AV suspends all jobs. In this case, this threshold has not yet been reached.

Solution: Contact Quorum Support to remove archive jobs as there is no user facility in the onQ Archive Vault to remove jobs, or contact Quorum Sales to discuss purchasing additional storage.

13.16.3 A0103

Message: The free Repository space on *ApplianceName* has fallen to critical limits. onQ is starting to expire the oldest backups for each Protected Node in order to recover space.

Cause: onQ has a built-in policy whereby if its repository disk space utilization exceeds 85% onQ begins to "free up" space by deleting PN backups, starting with the oldest backups. In this case, this threshold was reached, so onQ began deleting PN backups, starting with the oldest backups.

Solution: Nothing. The system is working as expected.

13.16.4 A0104

Message: The free Repository space on *ApplianceName* has fallen to critical limits. The Archive Vault is suspending all archive jobs until additional space becomes available.

Cause: AV has a built-in policy whereby if its repository disk space utilization exceeds 85%, AV suspends all jobs. In this case, this threshold was reached. Protection remains ON, but no further archiving will take place.

Solution: Contact Quorum Support to remove archive jobs as there is no user facility in the onQ Archive Vault to remove jobs, or contact Quorum Sales to discuss purchasing additional storage.

13.16.5 A0303

Message: The free Repository space on *ApplianceName* has fallen below allowable limits. Expiration of oldest backups failed for some PNs. onQ successfully expired *PercentageExpired* oldest backups: *NumberOfExpiredByPN*

Cause: onQ has a built-in policy whereby if its repository disk space utilization exceeds 85% onQ begins to “free up” space by deleting PN backups, starting with the oldest backups. In this case, this threshold was reached, so onQ tried deleting PN backups, starting with the oldest backups. Some expirations failed.

Solution: Contact Quorum Support to help with expirations, then Quorum Sales to add more disk space.

13.16.6 A0304

Message: The free Repository space on *ApplianceName* has fallen below allowable limits. onQ expired oldest *PercentageExpired* of backups (total of *NumberOfExpiredBackups*) for each Protected Node in order to recover space:" *NumberOfExpiredByPN*.

Cause: onQ has a built-in policy whereby if its repository disk space utilization exceeds 85% onQ begins to “free up” space by deleting PN backups, starting with the oldest backups. In this case, this threshold was reached, so onQ began deleting PN backups, starting with the oldest backups.

Solution: Delete orphan backups manually, reduce retention period, or remove a node for which you don't need protection. Contact Sales to discuss options for expanding disk space.

13.17 onQ Disk Space Alerts

The onQ Appliance needs a certain amount of disk space to do its job. Fortunately, there are ways to free up disk space. For example, you can change your backup retention policy.

For specific alerts:

- A0201
- A0301
- A0302

13.17.1 A0201

Message: The free OS disk space on the onQ appliance *ApplianceName* has fallen to critical limits and the automated cleanup process was unable to resolve the issue.

Cause: The onQ Appliance needs a certain amount of disk space to do its job.

Solution: Contact Quorum Support immediately.

13.17.2 A0301

Message: While in critically low free disk space state, current expiration policy allowed no backups to be expired on *ApplianceName* appliance. The condition will remain, please add physical disk space to the appliance.

Cause: The onQ Appliance needs a certain amount of disk space to do its job. onQ relies on the ability to expire backups in an effort to "free up" disk space. Your policy forbids expirations.

Solution: Change your policy to allow for expirations. Go to Change Backup Retention Policy.

13.17.3 A0302

Message: While in critically low free disk space state, attempt to expire oldest snapshots failed. Protection is being stopped.

Cause: The onQ Appliance needs a certain amount of disk space to do its job. onQ relies on the ability to expire backups in an effort to recover disk space; however, it was unable to do so.

Solution: Contact Quorum Support immediately.

13.18 Connection Alerts

onQ sends connection alerts when the DR system loses its connection with the HA Appliance, then re-establishes the connection. Depending on the circumstances, you might need to exercise your Disaster Recovery procedures.

For specific alerts:

- A0705
- A0706
- A0707
- A0708

13.18.1 A0705

Message: The onQ HA Manager *onQHostName* cannot be contacted. If this is confirmed to be a site failure, you will need to start Disaster Recovery protection from *ApplianceName*.

Cause: Possible causes are:

The HA appliance is completely down (unavailable) – including protection. However, DR is up and running.

The ‘Linkmon’ service on DR Appliance cannot start or isn’t working properly.

The onQ (Remote) configuration on DR Appliance does not have correct hostname or IP address for HA Appliance.

Solution: If this is a site failure, go to "(Workflow) Fail over HA Appliance to DR Appliance" on page 268.

13.18.2 A0706

Message: The onQ Manager *onQHostName* cannot be contacted.

Cause: Possible causes are:

- The DR Appliance is completely down (unavailable) – including protection. However, the onQ Archive Vault is up and running.
- The 'Linkmon' service on AV virtual machine cannot start or isn't working properly.

Solution: Restart the DR Appliance (see To restart the onQ Appliance:), if it isn't online. If the DR Appliance is up and running, restart protection on the AV: (1) From the drop-down menu > **Stop Archive Protection**, then (2) drop-down menu > **Start Archive Protection**.

13.18.3 A0707

Message: The link between the onQ Managers *onQHostName* and *ApplianceName* has been re-established. If any RNs on *ApplianceName* were being used in production mode, you will need to schedule your failback procedure before bringing up your PNs to minimize data loss.

Cause: Possible causes are:

- The HA appliance is completely down (unavailable), including protection. However, DR is up and running.
- The 'Linkmon' service on DR Appliance cannot start or isn't working properly.
- The onQ (Remote) configuration on DR Appliance does not have correct hostname or IP address for HA Appliance.

Solution: If this was a site failure, "(Workflow) Fail back DR to HA Appliance" on page 270.

13.18.4 A0708

Message: The link to the onQ Manager *onQHostName* has been re-established.

Cause: Disconnects can happen for a variety of reasons.

Solution: Nothing. This message is for informational purposes only. The HA Appliance is now accessible and the DR Appliance has re-established a connection with the HA Appliance.

13.19 Self-Test Alerts

Some self-test failures can easily be resolved by making sure that the RN has the correct network configuration. However, other failures are quite complicated because there are third party applications that can interfere with self-tests.

For specific alerts:

- A1001
- A1002
- A1003

13.19.1 A1001

Message: The automated self-test of *RNName* running on *ApplianceName* was successful. The snapshot level was *time timeZone mm*/*dd*/*yy*. No further notifications will be sent unless a self-test failure is detected.

Cause: The self-test started per the policy set in Configure automatic testing of RNs.

Solution: Nothing. Self-test is running as expected.

13.19.2 A1002

Message: The automated self-test of *RNName* running on *ApplianceName* failed. The failed snapshot level was *time timeZone mm*/*dd*/*yy*. No further notifications will be sent unless a self-test of this RN is successful.

Cause: onQ defines a failure as any self-test that doesn't complete successfully within a 15-minute period.

Solution: Perform the series of solutions outlined in Self-Test and RN Boot Problems. If these solutions don't work, contact Quorum Support.

13.19.3 A1003

Message: The automated self-test of *RNName* running on *ApplianceName* was terminated before completion due to resource constraints. Self-test will continue when sufficient resources are available.

Cause: Possible causes are:

- onQ doesn't have sufficient resources; as such, self-test termination is expected behavior. The self-test will work next time.
- The same or a different RN was brought up in production mode. Any running self-tests terminate immediately.

Solution: If there are insufficient resources, you don't need to do anything: the self-test will restart on its own. If due to another RN in production, the self-tests will not run until you turn off the RN in production mode.

13.20 Change iDRAC network configuration

iDRAC (Integrated Dell Remote Access Controller) is a card on the onQ Appliance. This hardware enables Quorum Support to observe the onQ Appliance boot. If your onQ Appliance model is a T20, you do not have iDRAC. All other models have either an iDRAC6 or iDRAC7.

To change iDRAC's network configuration:

This procedure is intended for Quorum Support only. Unless instructed to do so by Quorum Support, do not change iDRAC's network configuration.

1. Log on to the onQ Appliance's onQ Portal. You must log on as `varadmin` user.
2. Click the **APPLIANCE CONFIG** tab > **HYPERVISOR** page.
3. Click the **iDRAC** button. The iDRAC dialog appears.

   The onQ Portal displays an error if your onQ Appliance does not have iDRAC. If you know that the onQ Appliance has iDRAC, click the iDRAC button again: Sometimes a temporary networking glitch can cause iDRAC to become unavailable.
4. Click **MODIFY**, make your changes, then **SAVE**.

14

User Management

- Add users
- Delete users
- Change user passwords

14.1 Add users

By default, the onQ Appliance provides pre-defined user names (`VARAdmin`, `Admin`, `onQRestore`, `Operator`, and `archiveuser`.

However, you can add to this list if you have additional administrators that need access to the onQ Appliance; for auditing purposes, this approach is a best practice.

To add a user:

1. Log on to either the HA Appliance's onQ Portal or the DR Appliance's onQ Portal.
2. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **Users** page.

   You see the list of users at or below your login privilege level.
3. Click the plus button (**+**).
4. Specify username and password. User names are not case sensitive, but passwords are case sensitive.
5. Choose a user role:

   User roles provide the privileges for a user. By default, onQ provides pre-defined user roles:
   - **Administrator** - can add/modify PN configurations (i.e. change backup parameters), but cannot set up the onQ Appliance. In addtion, users with this role can upload packages.
   - **Monitor** - cannot make changes (read-only).
   - **Operator** - cannot make changes (read-only), but can start RNs.

- **Restore** - can set up a share. This role is used by onQ only; you cannot assign users this role.
- **VARAdmin** - can do anything, the only one that can modify/reconfigure the onQ Appliance (software, settings). VAR is short for Value-Added Reseller. If you use Hybrid Cloud, you cannot assign users this role. This user cannot be deleted.

6. Specify a timeout value, then **SAVE**.

   onQ automatically logs off this user after this many seconds of inactivity. For security purposes, don't set `Admin` or `VARAdmin` to 0, which means never automatically log off. Ideally, retain the default of 300 seconds.

Related Topics

Delete users

14.2 Delete users

You cannot delete the pre-defined users named *archiveuser* or *VARAdmin*.

To delete a user:

1. Log on to either the HA Appliance's onQ Portal or the DR Appliance's onQ Portal.
2. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **Users** page.
3. Select a user, then click the minus button (**-**).
4. Click **Yes** to confirm.

Related Topics

Add users

14.3 Change user passwords

For security purposes, the default password for your Appliance(s) cannot be made public. If you do not know your default password, contact Quorum Support.

When you first configured your Appliance(s), you logged on as `VARAdmin` or, in the case of Hybrid Cloud, `Admin` using the default password.

After you configure your Appliance(s) and, if applicable, the onQ Archive Vault, you must immediately change the default passwords for these users. The user ID is not case-sensitive, but the password is case sensitive.

Note: It is critical that you remember your VARAdmin password: Quorum Support cannot recover it and cannot perform certain support tasks without it. VARAdmin is not needed for day-to-day administration and operation. This caveat does not apply to Hybrid Cloud users as you do not have access to VARAdmin credentials.

To change a user password:

1. Log on to the HA Appliance's onQ Portal or the DR Appliance's onQ Portal as `VARAdmin`, or `Admin`. Each Appliance has its own password.
2. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **Users** page.
3. Select a user, then **MODIFY**.
4. Specify a new username and password.

   Best practice is to use a strong password.

15

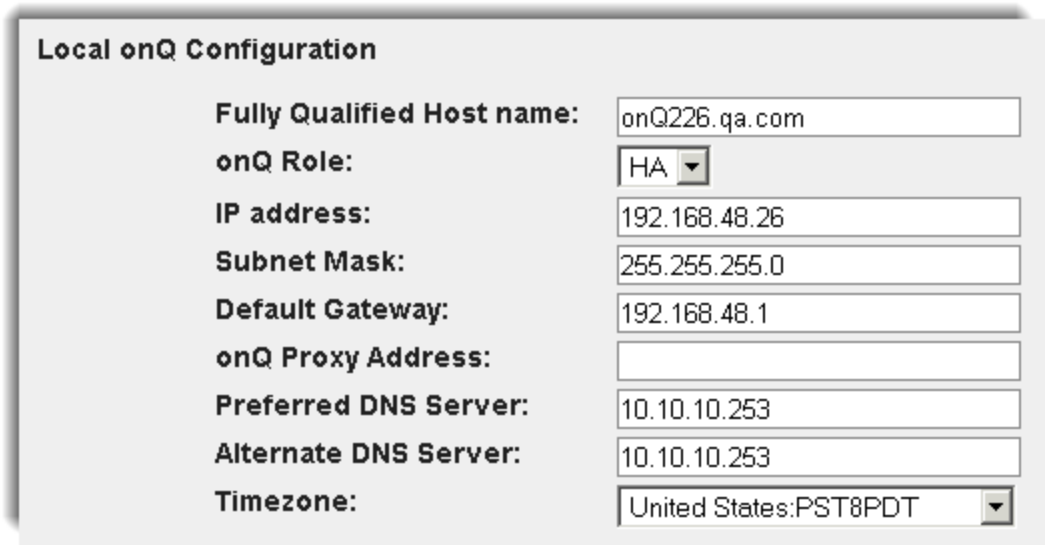
Example Configurations

- Local: Example of HA Appliance
- Local: Example of DR Appliance
- Local: Example of DR Mirror
- Remote: Example of HA Appliance with Hybrid Cloud
- Remote: Example of HA Appliance with Remote DR Appliance
- Remote: Example of DR Appliance with DR Mirror
- Remote: Example of DR Mirror

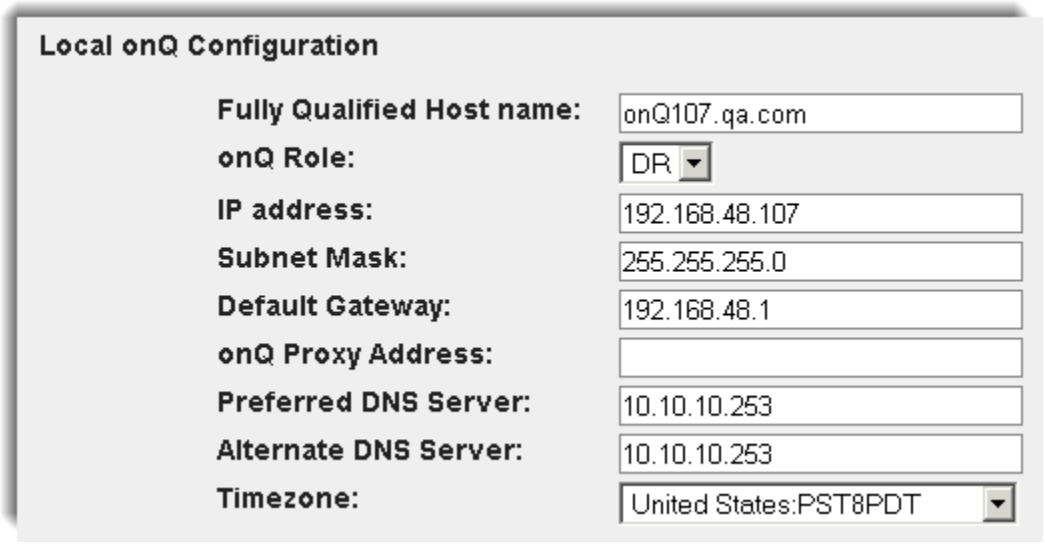
15.1 Local onQ Appliance Configurations

- Local: Example of DR Appliance
- Local: Example of DR Mirror

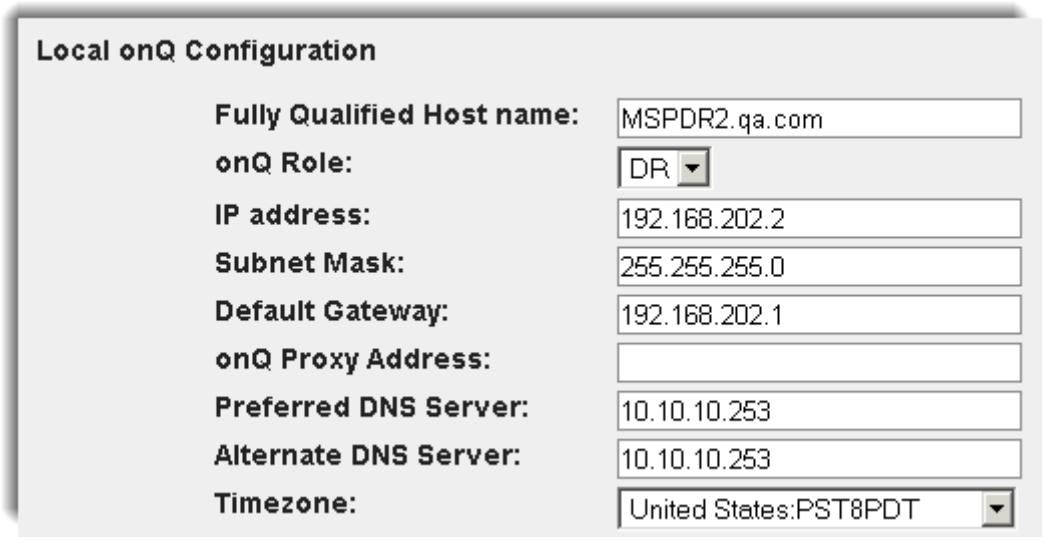
15.1.1 Local: Example of HA Appliance



15.1.2 Local: Example of DR Appliance



15.1.3 Local: Example of DR Mirror



15.2 Remote onQ Appliance Configurations

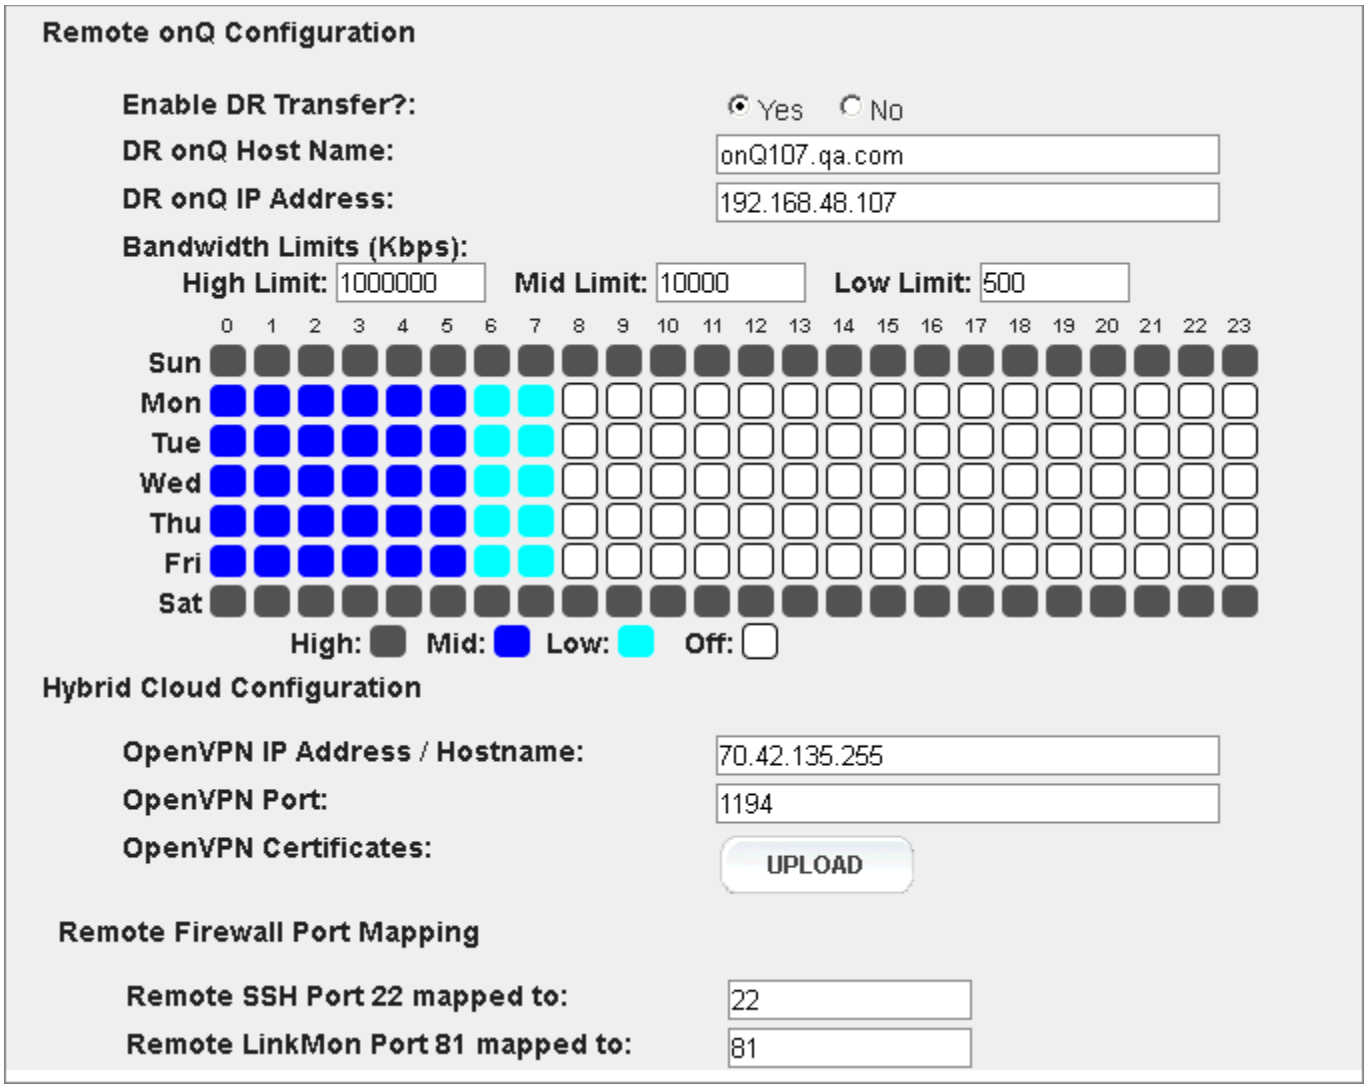
- Remote: Example of HA Appliance with Hybrid Cloud

- Remote: Example of HA Appliance with Remote DR Appliance
- Remote: Example of DR Appliance with DR Mirror
- Remote: Example of DR Mirror

15.2.1 Remote: Example of HA Appliance with Hybrid Cloud

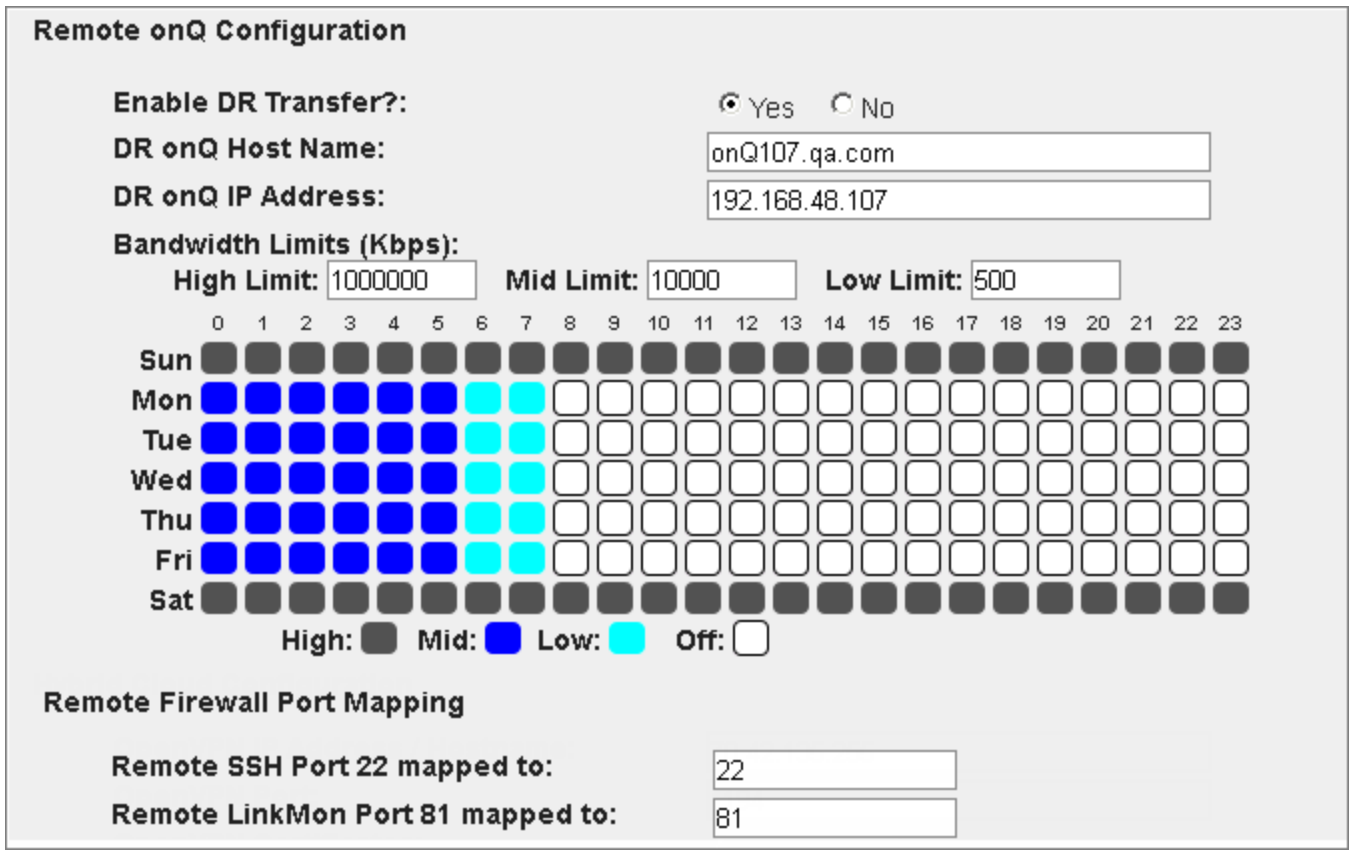
In order for the HA Appliance to communicate with the DR Appliance, via a secure VPN connection, you must provide the following information for the remote Appliance.

- DR onQ IP Address/Hostname - Quorum-assigned private IP address. If you do not know your IP address, contact Quorum Support.
- OpenVPN IP Address/Hostname - Quorum-assigned public IP address so that the HA Appliance can VPN to the DR Appliance. If you do not know your IP address, contact Quorum Support.
- DR OpenVPN Port - 1194
- OpenVPN Certificates - Quorum-provided certificates. In some cases, Quorum Support might upload these for you in advance.


Remote onQ Configuration
Enable DR Transfer?:
Yes
No
DR onQ Host Name:
onQ107.qa.com
DR onQ IP Address:
192.168.48.107
Bandwidth Limits (Kbps):
High Limit: 1000000
Mid Limit: 10000
Low Limit: 500
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23
Sun
Mon
Tue
Wed
Thu
Fri
Sat
High:
Mid:
Low:
Off:
Hybrid Cloud Configuration
OpenVPN IP Address / Hostname:
70.42.135.255
OpenVPN Port:
1194
OpenVPN Certificates:
UPLOAD
Remote Firewall Port Mapping
Remote SSH Port 22 mapped to:
22
Remote LinkMon Port 81 mapped to:
81

15.2.2 Remote: Example of HA Appliance with Remote DR Appliance

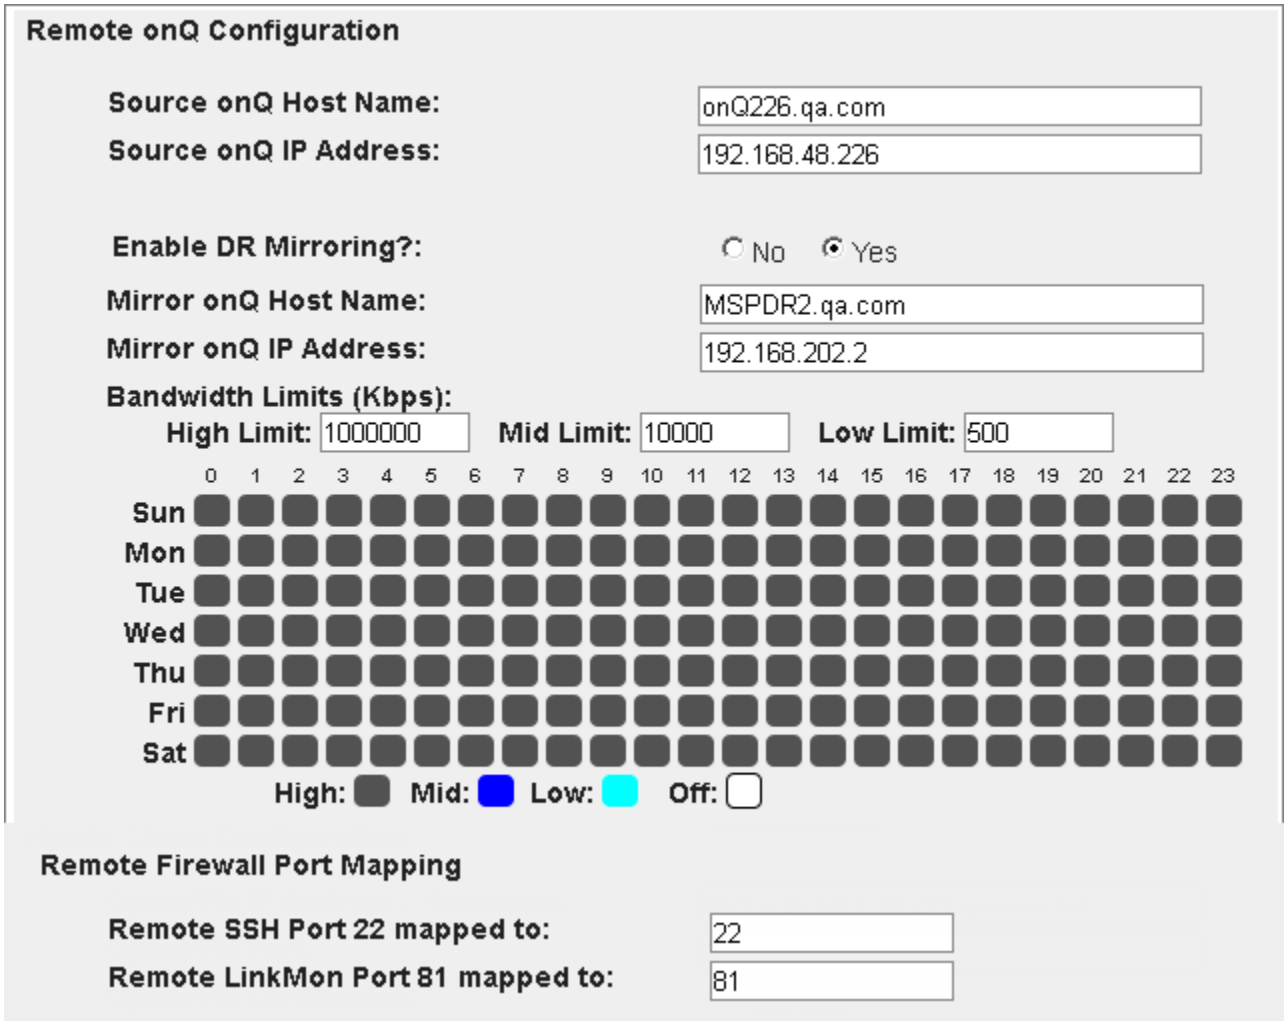
Warning: Enable DR Transfer?. After an initial configuration, do not disable DR transfers as outlined in Disable replication globally. If you do so, your repositories will be out of sync.



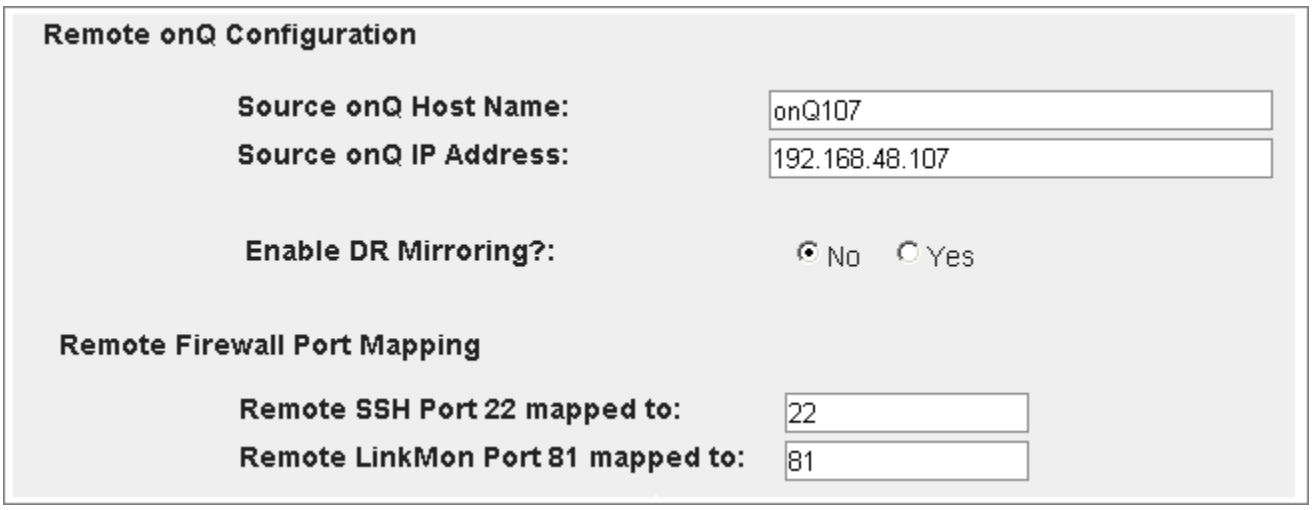
15.2.3 Remote: Example of DR Appliance with DR Mirror

Warning: Enable DR Mirroring? After an initial configuration, do not disable DR mirroring as outlined in Enable and disable DR Mirroring. If you do so, your repositories will be out of sync.

Warning: Do not use HA Appliance information in place of DR mirroring information. Doing so creates an endless loop between HA Appliance and DR Appliance.


Remote onQ Configuration
Source onQ Host Name:
onQ226.qa.com
Source onQ IP Address:
192.168.48.226
Enable DR Mirroring?:
No
Yes
Mirror onQ Host Name:
MSPDR2.qa.com
Mirror onQ IP Address:
192.168.202.2
Bandwidth Limits (Kbps):
High Limit: 1000000
Mid Limit: 10000
Low Limit: 500
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23
Sun
Mon
Tue
Wed
Thu
Fri
Sat
High:
Mid:
Low:
Off:
Remote Firewall Port Mapping
Remote SSH Port 22 mapped to:
22
Remote LinkMon Port 81 mapped to:
81

15.2.4 Remote: Example of DR Mirror



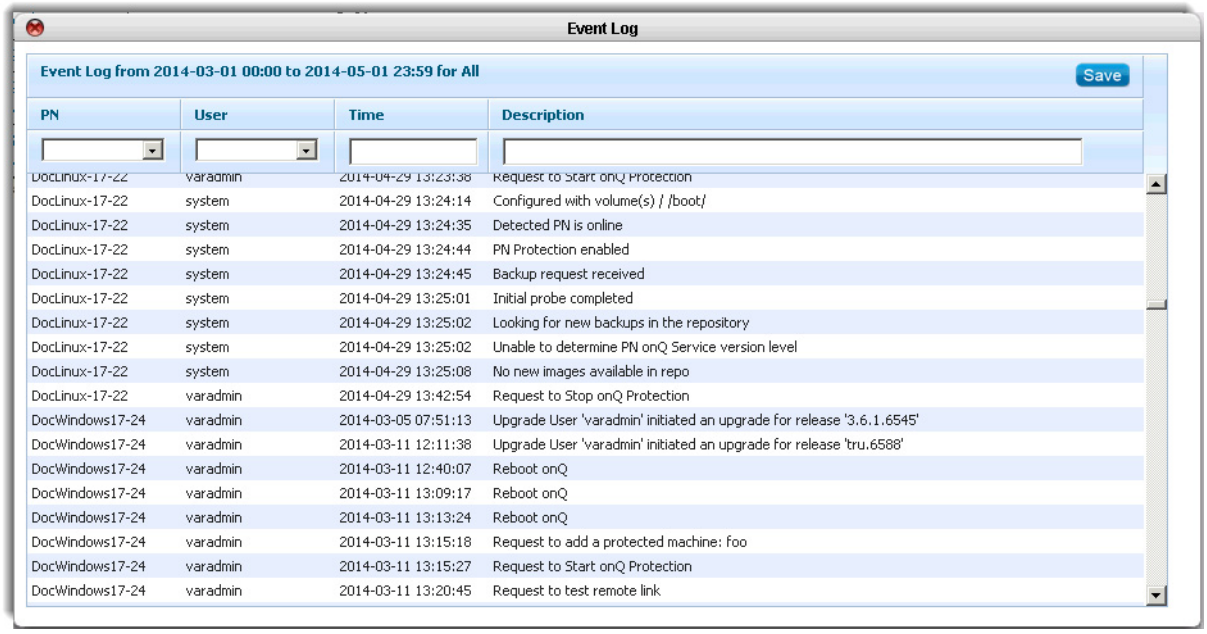
16

Example Logs

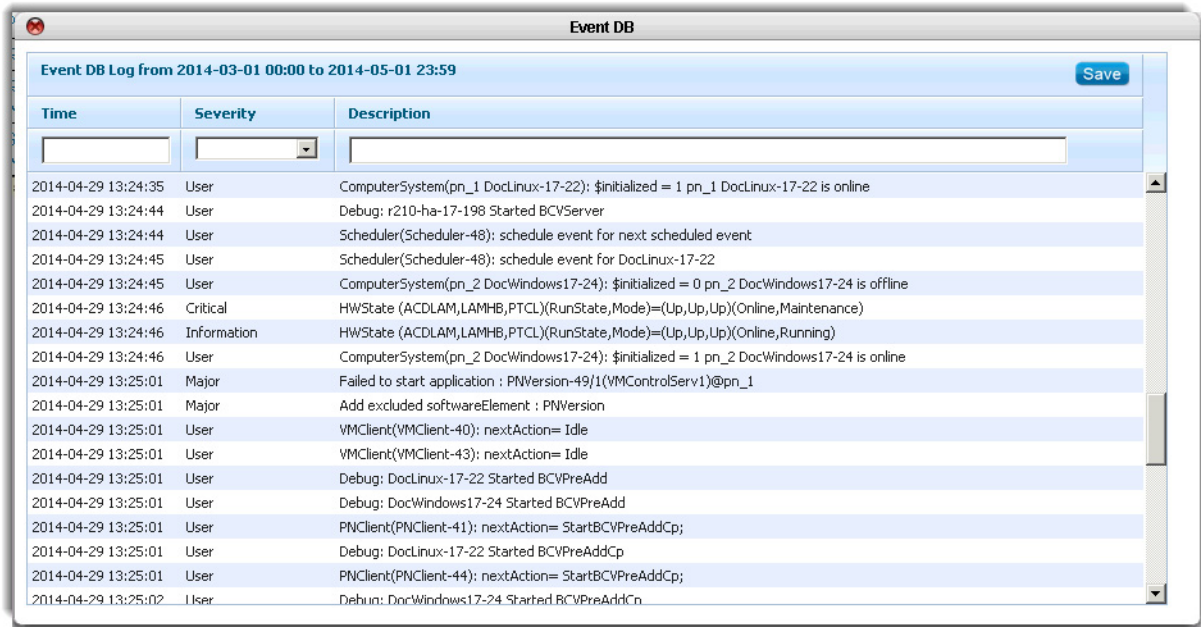
- Example Event Log
- Example Event DB Log
- Example Expired Snapshot Log
- Example FLR Activity Log
- Example HA -> DR Transfer Log
- Example Manager Debug Log
- Example PN Configuration Log
- Example Self Test Log
- Example Upgrade Log
- Example WSR Activity Log

The following screen shots represent example logs.

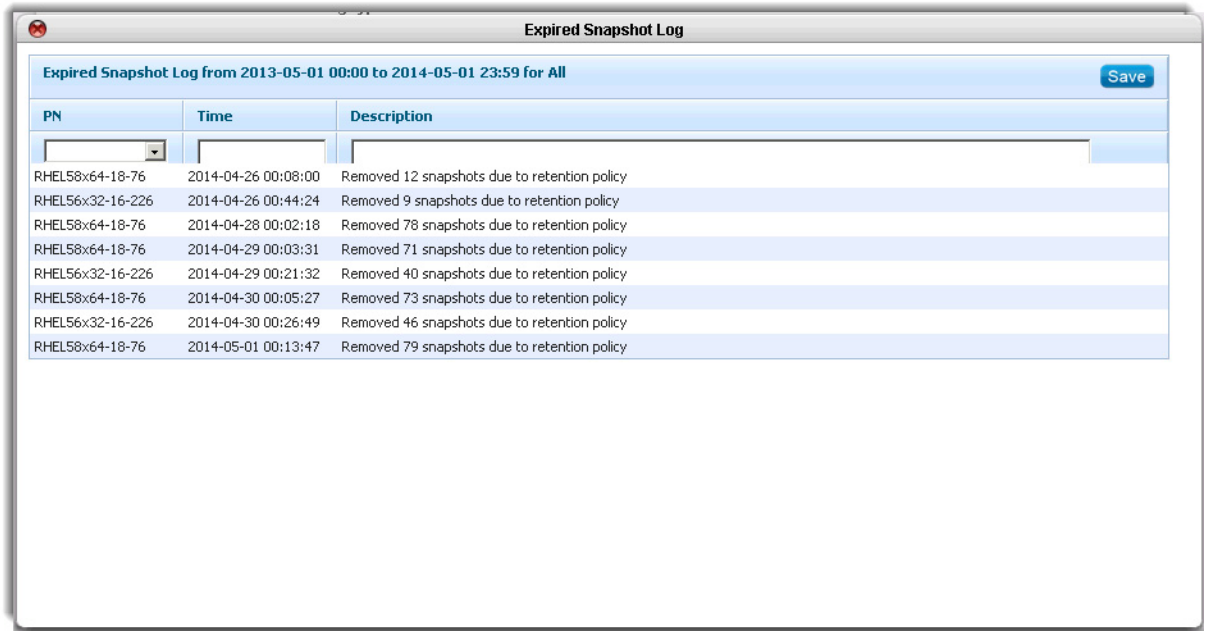
16.1 Example Event Log



16.2 Example Event DB Log



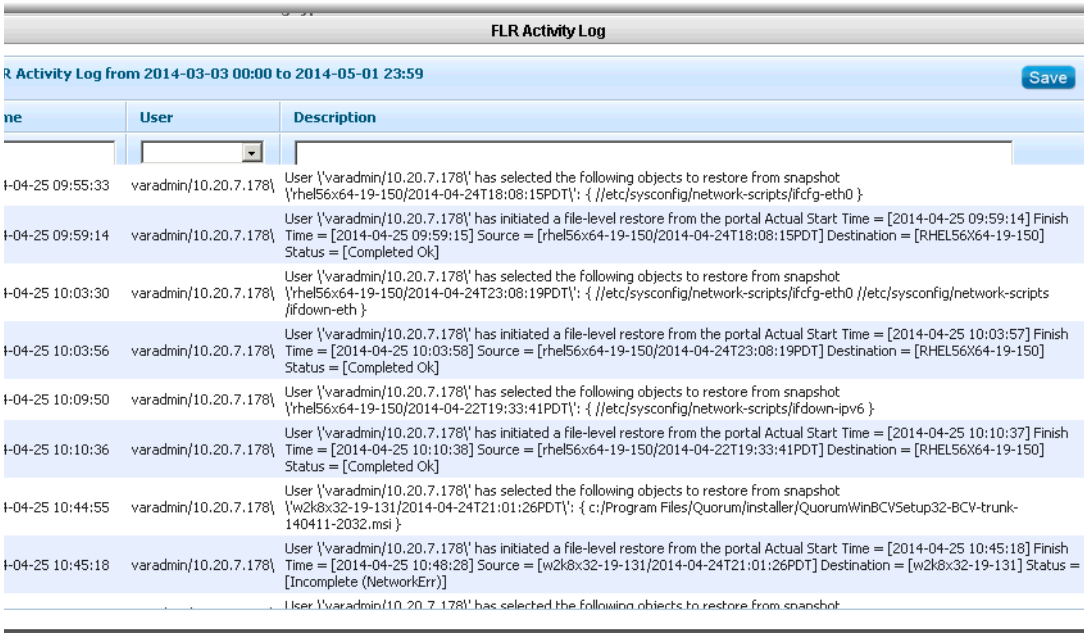
16.3 Example Expired Snapshot Log



16.4 Example FLR Activity Log

If you'd like to see a graphical representation of this data, consider using the

Restore Detail Report outlined in Generate on-demand reports.

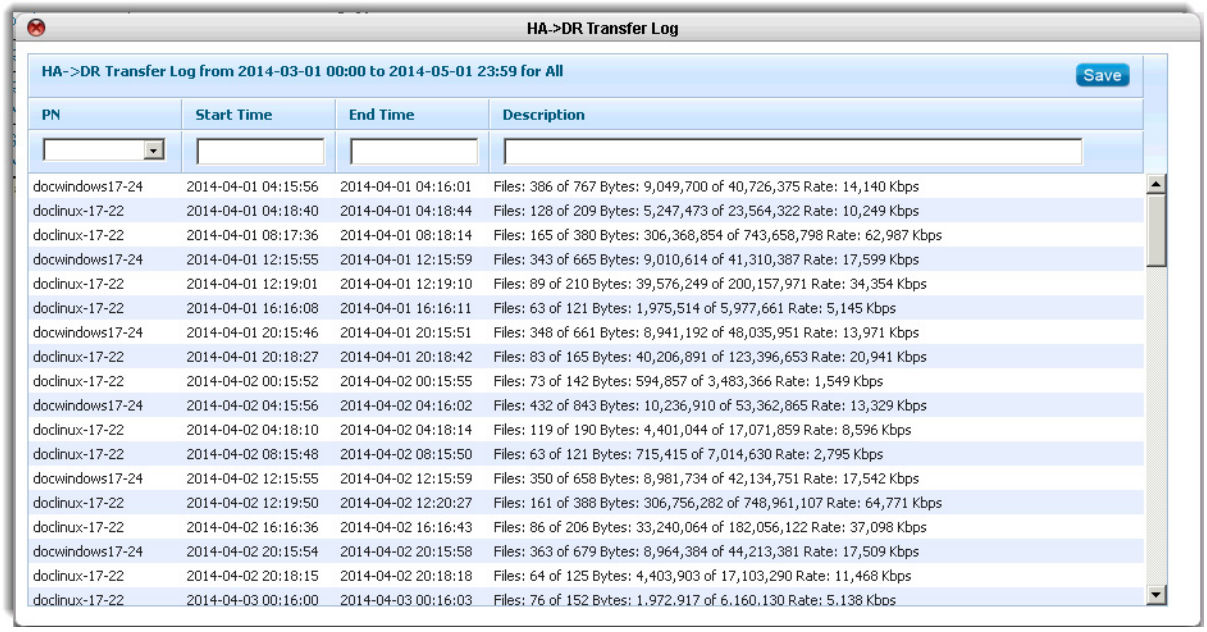


FLR Activity Log

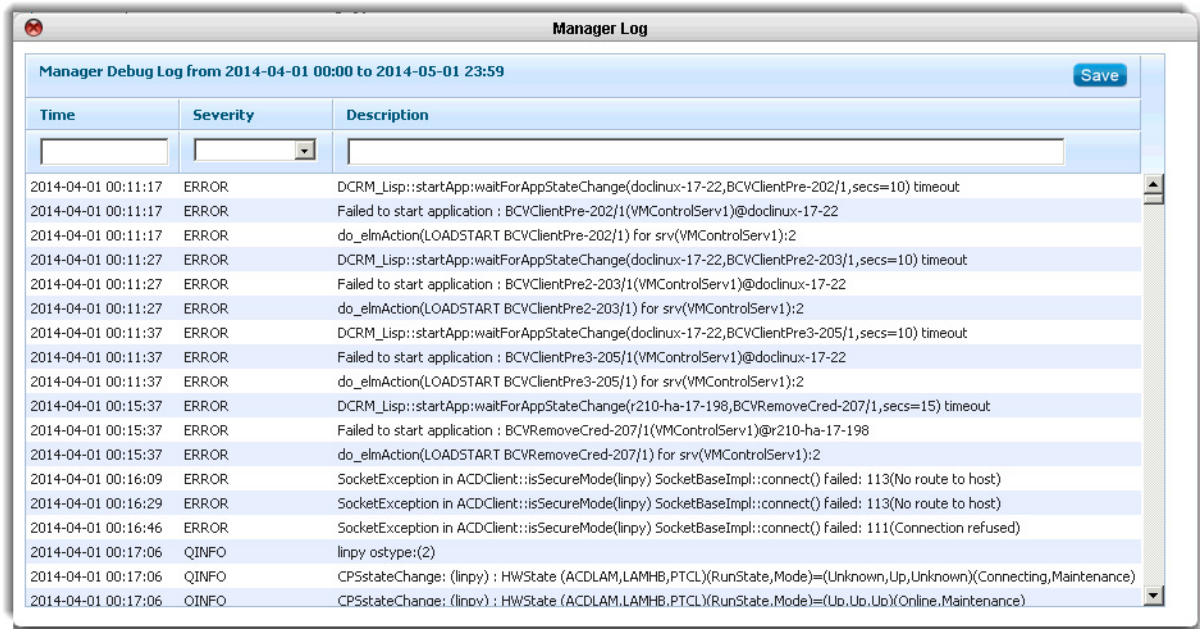
R Activity Log from 2014-03-03 00:00 to 2014-05-01 23:59 Save

ne	User	Description
4-04-25 09:55:33	varadmin/10.20.7.178\	User \'varadmin/10.20.7.178\' has selected the following objects to restore from snapshot \'rhel56x64-19-150/2014-04-24T18:08:15PDT\': { //etc/sysconfig/network-scripts/ifcfg-eth0 }
4-04-25 09:59:14	varadmin/10.20.7.178\	User \'varadmin/10.20.7.178\' has initiated a file-level restore from the portal Actual Start Time = [2014-04-25 09:59:14] Finish Time = [2014-04-25 09:59:15] Source = [rhel56x64-19-150/2014-04-24T18:08:15PDT] Destination = [RHEL56X64-19-150] Status = [Completed Ok]
4-04-25 10:03:30	varadmin/10.20.7.178\	User \'varadmin/10.20.7.178\' has selected the following objects to restore from snapshot \'rhel56x64-19-150/2014-04-24T23:08:19PDT\': { //etc/sysconfig/network-scripts/ifcfg-eth0 //etc/sysconfig/network-scripts /ifdown-eth }
4-04-25 10:03:56	varadmin/10.20.7.178\	User \'varadmin/10.20.7.178\' has initiated a file-level restore from the portal Actual Start Time = [2014-04-25 10:03:57] Finish Time = [2014-04-25 10:03:58] Source = [rhel56x64-19-150/2014-04-24T23:08:19PDT] Destination = [RHEL56X64-19-150] Status = [Completed Ok]
4-04-25 10:09:50	varadmin/10.20.7.178\	User \'varadmin/10.20.7.178\' has selected the following objects to restore from snapshot \'rhel56x64-19-150/2014-04-22T19:33:41PDT\': { //etc/sysconfig/network-scripts/ifdown-ipv6 }
4-04-25 10:10:36	varadmin/10.20.7.178\	User \'varadmin/10.20.7.178\' has initiated a file-level restore from the portal Actual Start Time = [2014-04-25 10:10:37] Finish Time = [2014-04-25 10:10:38] Source = [rhel56x64-19-150/2014-04-22T19:33:41PDT] Destination = [RHEL56X64-19-150] Status = [Completed Ok]
4-04-25 10:44:55	varadmin/10.20.7.178\	User \'varadmin/10.20.7.178\' has selected the following objects to restore from snapshot \'w2k8x32-19-131/2014-04-24T21:01:26PDT\': { c:/Program Files/Quorum/installer/QuorumWinBCVSetup32-BCV-trunk-140411-2032.msi }
4-04-25 10:45:18	varadmin/10.20.7.178\	User \'varadmin/10.20.7.178\' has initiated a file-level restore from the portal Actual Start Time = [2014-04-25 10:45:18] Finish Time = [2014-04-25 10:48:28] Source = [w2k8x32-19-131/2014-04-24T21:01:26PDT] Destination = [w2k8x32-19-131] Status = [Incomplete (NetworkErr)]
		User \'varadmin/10.20.7.178\' has selected the following objects to restore from snapshot

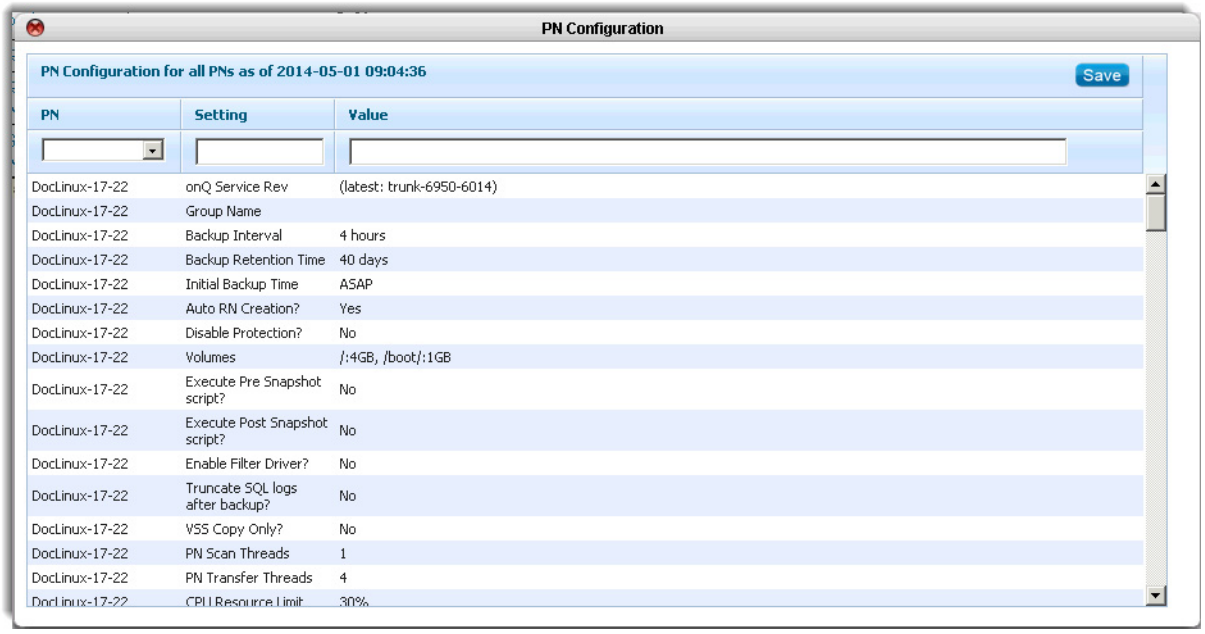
16.5 Example HA -> DR Transfer Log



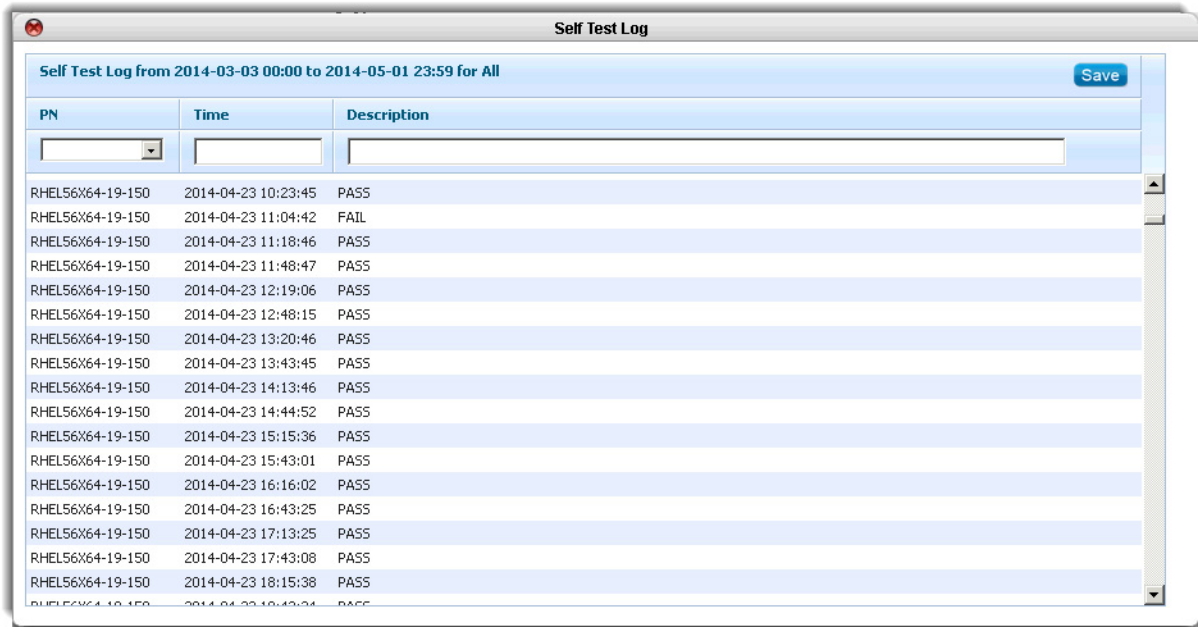
16.6 Example Manager Debug Log



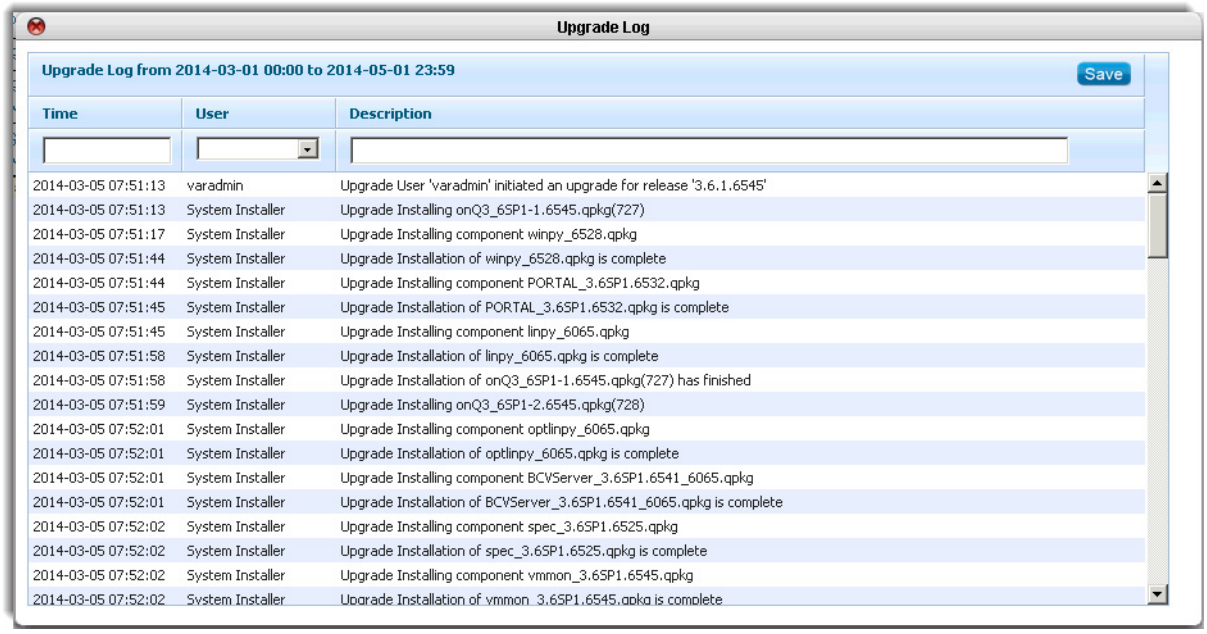
16.7 Example PN Configuration Log



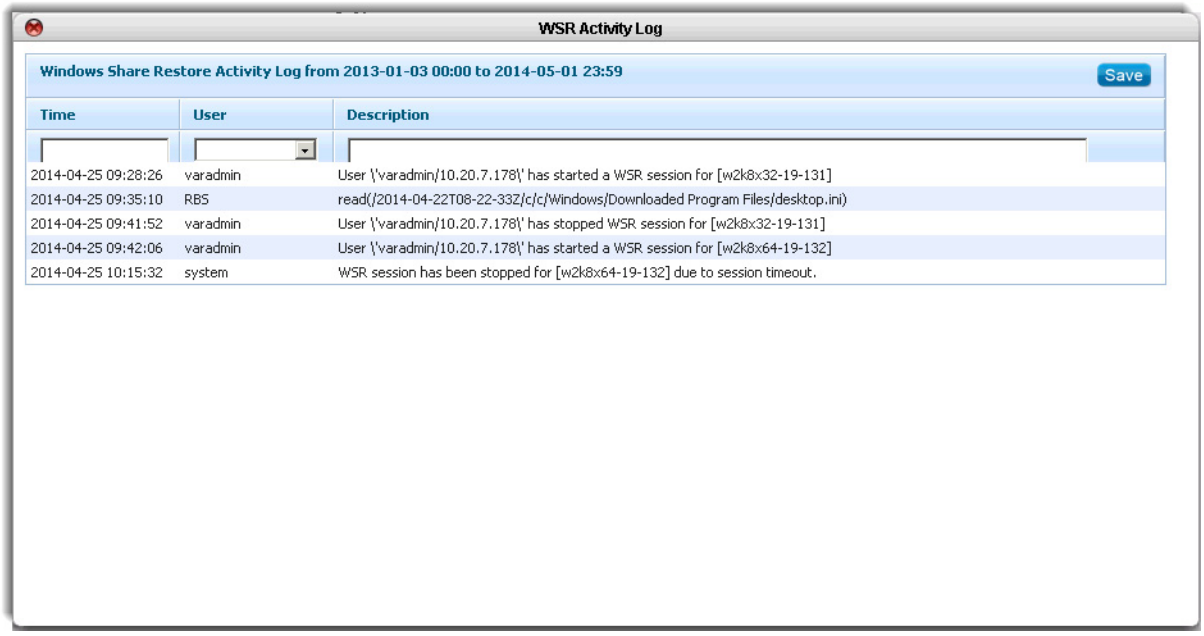
16.8 Example Self Test Log



16.9 Example Upgrade Log



16.10 Example WSR Activity Log



17

Tech Notes

- Install kernel-xen RPM package

17.1 Install kernel-xen RPM package

RNs that are running Redhat 5.x depend on a kernel-xen RPM package. If you are running Redhat 6.x this package is installed by default; as such, you don't need to perform the following procedure.

Perform this procedure if you enrolled your PNs without having previously installed the kernel-xen RPM package, or if you're attempting to restore a Recovery PN as part of a BMR. Redhat 5.x RNs will not boot without this package.

To install the kernel-xen RPM package:

The following procedure works on `kernel-xen-2.6.18`, and should work for all subsequent revisions.

1. Determine if you need the kernel-xen RPM package and, if so, which package:

   a. Verify the OS version. Again, the PN must be Redhat 5.x version.

```
# cat /etc/redhat-release
Red Hat Enterprise Linux Server release 5.7 (Ti-
kanga)
```

   b. Determine if the PN's OS is 32-bit or 64-bit.

```
# uname -a
Linux RHEL57x64-18-178 2.6.18-274.el5 #1 SMP Fri
Jul 8 17:36:59 EDT 2011 x86_64 x86_64 x86_64
GNU/Linux
```

In this case, `x86_64` means the 64-bit package is required.

2. Download the kernel-xen package for RedHat 5.x PN (x32 or x64).
   - For 32-bit, go to http://mirror.centos.org/centos-5/5/os/i386/CentOS/, scroll down to view the complete package name for the latest `kernel-xen` package, then run the following command:

```
# wget
http://mirror.centos.org/centos/5/os/i386/CentOS
/kernel-xen-<latestRelease>-
<latestRevision>.el5.i686.rpm
```

   - For 64-bit, go to http://mirror.centos.org/centos-5/5/os/x86_64/CentOS/, scroll down to view the complete package name for the latest `kernel-xen` package, then run the following command:

```
# wget
http://mirror.centos.org/centos/5/os/x86_64/Cent
OS/kernel-xen-<latestRelease>-
<latestRevision>.el5.x86_64.rpm
```

3. Install the kernel-xen package.
   - For 32-bit:

```
# rpm -iv kernel-xen-<latestRelease>-
<latestRevision>.el5.i686.rpm
```

   - For 64-bit:

```
# rpm -iv kernel-xen-<latestRelease>-
<latestRevision>.el5.x86_64.rpm
```

4. Make the `initrd` file, which will be used to build the RN.

```
#mkinitrd -v -f --preload=xennet --preload=xenblk
--omit-scsi-modules --omit-raid-modules \
/boot/initrd-<latestRelease>-
<latestRevision>.el5xen.img.5 <latestRelease>-
<latestRevision>.el5xen
```

5. Make the `/boot/grub/grub.conf.xvf5` file, which signals the RN to use it as grub boot menu.

The /boot/grub/grub.conf.xvf and /boot/grub/grub.conf.xvf[1-4] must not be installed because they have higher priority than /boot/grub/grub.conf.xvf5.

```
# vi /boot/grub/grub.conf.xvf5
```

6. Cut and paste the following line to the file.

```
# grub.conf generated by Quorum onQ

# The root device path is not important and will
be changed automatically during RN build.

default=0

timeout=5

hiddenmenu

title Red Hat Enterprise Linux Server by Quorum
onQ (<latestRelease>-<latestRevision>.el5xen)

        root (hd0,0)

        kernel /vmlinuz-<latestRelease>-
<latestRevision>.el5xen ro root=/dev/xvda1
rd_NO_LUKS rd_NO_MD rhgb crashkernel=auto
rd_NO_LVM
       initrd /initrd-<latestRelease>-<latestRev-
ision>.el5xen.img.5
```

If you're performing this procedure on the RN (not the PN), back up and build the RN.

To uninstall the kernel-xen RMP package:

1. Uninstall the kernel-xen rpm:

```
# rpm -e kernel-xen-<latestRelease>-
<latestRevision>.el5
```

2. Remove the /boot/grub/grub.conf.xvf5 file.

```
# \rm -f /boot/grub/grub.conf.xvf5
```

18

Glossary

onQ Appliance ..... A separate hardware component designed to host a specific computing resource. onQ Appliance is a separate computer system hosting a virtualization hypervisor, OS, Quorum-specific monitoring and other software, and sufficient resources to be able to support intermediate backups and when recovery is required, able to run a large number of nodes as virtual machines.

business continuity The totality of policies and actions put in place by a business in order to maintain its operations and relationships with partners and clients even when interruptions or disasters occur.

de-duplication ..... The elimination of redundant blocks of data across PNs to reduce storage and bandwidth requirements. De-duplication is important in backup systems where timely recovery is a concern.

disaster recovery .. The ability to recover business systems at an off-site facility in the event of a primary site failure.

DR Appliance ..... The DR Appliance provides disaster recovery protection for your nodes. Often the DR Appliance is truly remote, being located in a different geographical location and connected through a separate WAN, but a more important distinction is the role it plays as the provider of disaster recovery protection.

HA Appliance ..... The HA Appliance is the one connected directly to the network on which the Protected Nodes are located. The HA Appliance provides high availability protection.

high availability..... Protection of nodes by onQ located on the same LAN as the nodes it protects. When HA recovery is required, Recovery Nodes are started on the HA Appliance.

hypervisor ..... hypervisor is software platform that allows multiple oper-

ating systems to run concurrently as virtual machines on a host computer.

incremental backups onQ uses a backup scheme wherein a full backup is followed by a series of incremental backups, meaning backups that consist only of changes since the preceding backup. A Recovery Node can be quickly updated by simply merging the latest changes.

onQ Appliance ..... onQ Appliance refers to the hardware on which onQ Manager runs.

onQ Central ..... onQ Central is Quorum's support management system, onQ Central manages alert notifications, licenses, and updates.

onQ Manager ..... onQ Manager is the service that runs on an onQ Appliance coordinating all protection activities. onQ Manager sends alerts to onQ Central for processing.

onQ Service ..... onQ Service is the secure, lightweight software agent running on an agent-based Protected Node. It communicates with the onQ Manager.

onQ Monitor ..... onQ Monitor is a tool designed to enable you to monitor multiple onQ Appliances from one user interface.

onQ Portal ..... onQ Portal is the Web-based user interface used for all monitoring and management functions.

onQ virtual machine onQ virtual machine (also known as the *onQ instance)* is a separate operating system implemented in software. An onQ Appliance can host more than one onQ virtual machine. Each onQ instance provides an onQ Portal. An MSP (Managed Service Provider) configuration hosted by a single onQ Appliance can have several onQ instance

protection ..... Any node on your network for which onQ provides a virtual image that can be run is said to be protected by onQ.

Protected Node ..... The physical or virtual machine to be protected by onQ. Abbreviated **PN**.

Recovery Node ..... The up-to-date ready-to-run copy of the Protected Node.

Abbreviated **RN**.

replication ..... The process of reproducing a Recovery Node where the original image is on an HA Appliance and is replicated on the DR Appliance.

repository ..... The deduplicated archive of all snapshots of all of your Protected Nodes. From the repository (aka *snapshot repository*), onQ can reproduce a complete system image from any snapshot.

squirtcopy ..... onQ's backup utility/client. squirtcopy performs backups of the PNs to squirtserver. squirtcopy is packaged with linpy/winpy.

test mode..... Test mode means that a Recovery Node runs in a private network environment isolated from production facilities.

throttling..... Bandwidth throttling is a technique for limiting the amount of transfer bandwidth available for data transfer on a network.