Quorum®

onQ Administrator's Guide

Version 3.6 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 © 2008–2013 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 7

1.5 About Hybrid Cloud 9

1.6 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 14

2.3 Host-based Firewall and Anti-Virus Requirements 15

2.4 Operational Considerations 16

2.5 Identify interdependencies 16

3. Installation 17

3.1 Installation Overview 17

3.2 (Workflow) Install onQ Appliances 19

3.2.1 Step 1: Prepare to install 20

3.2.2 Step 2: Configure the onQ Appliances 20

3.2.3 Step 3: (Hybrid Cloud Only) Create VPN connection to Hybrid Cloud 22

3.2.4 Step 4: Connect the Appliance to your network 32

3.2.5 Step 5: Enroll protected nodes 33
3.2.6 Step 6: Start automated protection process 33
3.3 Initial Backup 34
3.4 Initialize the DR Repository 35

4. Updates 37
4.1 Upgrade Support 37
4.2 Update Appliance software 38
4.3 (Windows) Update node software 40
4.4 (Linux) Update node software 43
4.5 Get software packages 46
4.6 Delete software packages 48
4.7 Download software node packages 48
4.8 Install onQ Appliance license 49
4.9 View onQ software versions 51

5. The Basics 53
5.1 Log on to Appliance's user interface 54
5.2 Establish a VPN connection to Hybrid Cloud 55
5.3 Log off from Appliance's user interface 55
5.4 Pull-down menu 56
5.5 Reboot onQ Manager 58
5.6 Reboot Appliance 58
5.7 Shut down Appliance 59
5.8 Modify Appliance's network settings 59
5.9 Configure Appliance's hypervisor settings 60
5.10 Unlock UI 61
5.11 Add protected nodes manually 62
5.12 Modify protected nodes 64
5.13 Delete protected nodes 65
5.14 Protected Node Parameters 66

5.15 (Windows) Enroll protected nodes 69
5.16 (Linux) Enroll protected nodes 73

6. Protection 77
6.1 Start node protection 78
6.2 Stop node protection 78
6.3 Modify recovery node creation policy 79
6.4 Recreate recovery nodes 81
6.5 Disable protection of nodes 82
6.6 Create on-demand RNs 83
6.7 Delete on-demand RNs 84
6.8 Cancel RN updates 85
6.9 Manage hosts 86
6.10 Resize protected node's vdisk 87
6.11 Disable snapshot transfers 88
6.12 Enable and disable DR Mirroring 89
6.13 Enroll and Disenroll onQ Archive Vault 91

7. Backup and Restore 93
7.1 Backup and Recovery Workflow 94
7.2 Coordination of Backups and Recovery Nodes 96
7.3 Edit RN services list 99
7.4 Edit backup exclude list 101
7.5 Edit backup include list 104
7.6 Initiate immediate backups 105
7.7 Export or import Repository 106
7.8 Delete snapshots 108
7.9 Stop in-progress backups 108
7.10 Retrieve objects from Repository 109
7.10.1 Perform WSR Restore 110
7.10.2 Perform File-level Restore 113

7.11 Adjust bandwidth throttling 119
7.12 Limit resource utilization on PNs 121
7.13 Run custom backup scripts 121
7.14 Backup Activity Report 124
7.15 Methods of Performing Incremental Backups 126
7.16 Perform full rescan on PN 128
7.17 Back up and restore Oracle 11g database 129

8. **Bare Metal Restore** 133
   8.1 Perform Bare Metal Restore 134
   8.2 BMR Errors 141

9. **Security and Communications** 143
   9.1 Test HA-to-DR Link 144
   9.2 Set up trust relationships 145
   9.3 Modify TCP Ports for HA-DR Communications 151
   9.4 Synchronize system time 151
   9.5 Secure HA connection to PNs 152
   9.6 Restart the onQ Service 154

10. **Disaster Recovery and DR Preparedness** 155
    10.1 Fail over HA Appliance to DR Appliance 156
    10.2 Fail back DR to HA Appliance 158
    10.3 Change Appliance's role 160
    10.4 Start recovery nodes on HA Appliance 160
    10.5 Start recovery nodes on DR Appliance or DR Mirror 165
    10.6 Launch recovery node's console 167
    10.7 Test Mode 167
    10.8 Configure automatic testing of RNs 168
    10.9 Suspend and resume RN tests 169

11. **Monitoring** 171

11.1 Monitor multiple onQ Appliances 172
11.2 Monitor protected nodes 175
11.3 Monitor recovery nodes 179
11.4 Monitor backups 181
11.5 Monitor DR Appliance 186
11.6 Monitor onQ Archive Vault enrollment 190
11.7 Monitor disk space usage 191
11.8 Interpret e-mail alerts 192
11.8.1 License Expiration Alerts 193
11.8.2 Reboot Alerts 193
11.8.3 Protection Alerts 195
11.8.4 vDisk or Disk Space Alerts 196
11.8.5 Connection Alerts 198
11.9 Modify e-mail alert settings 199

12. Troubleshooting 201
12.1 Troubleshoot Installation Problems 202
12.2 Troubleshoot Connection Problems 203

13. Windows Cluster Services 205
13.1 Windows Cluster Services Overview 205
13.2 Download iSCSI initiator for Windows 2003 205
13.3 Configure Quorum onQ Service 206
13.4 Work with RNs in Windows Cluster Services Environment 212

14. User Management 221
14.1 Add users 221
14.2 Delete users 222
14.3 Change user passwords 223

15. Example Configurations 225
15.1 Local onQ Appliance Configurations 226
15.1.1 Local: Example of HA Appliance 226

15.1.2 Local: Example of DR Appliance 226
15.1.3 Local: Example of DR Mirror 227
15.2 Remote onQ Appliance Configurations 227
15.2.1 Remote: Example of HA Appliance with Hybrid Cloud 228
15.2.2 Remote: Example of HA Appliance with Remote DR Appliance 230
15.2.3 Remote: Example of DR Appliance with DR Mirror 231
15.2.4 Remote: Example of DR Mirror 233

16. Tech Notes 1
16.1 Search for EFI Partition 1
16.2 Make EFI Partition Bootable 4
16.3 Install kernel-xen RPM package 7

17. Glossary 11

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/3.6HF1/onQ/wwhelp/wwhimpl/js/html/wwhelp.htm#href=Getting%20Started/Welcome.html
AV	http://www.quorum.net/help/3.6HF1/Vault/wwhelp/wwhimpl/js/html/wwhelp.htm#href=Getting%20Started/Welcome.html

Chapter Organization

The chapters and appendices of this manual are as follows:

- Chapter 1, *Product Overview*
- Chapter 2, *Requirements*
- Chapter 3, *Installation*
- Chapter 4, *Updates*
- Chapter 5, *The Basics*
- Chapter 6, *Protection*
- Chapter 7, *Backup and Restore*
- Chapter 8, *Bare Metal Restore*
- Chapter 9, *Security and Communications*
- Chapter 10, *Disaster Recovery and DR Preparedness*
- Chapter 11, *Monitoring*

- Chapter 12, *Troubleshooting*
- Chapter 13, *Windows Cluster Services*
- Chapter 14, *User Management*
- Chapter 15, *Example Configurations*
- Chapter 16, *Tech Notes*
- Chapter 17, *Glossary*.

1

Product Overview

- About onQ
- Terminology
- How does onQ work?
- Inside onQ Appliance
- About Hybrid Cloud
- 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.

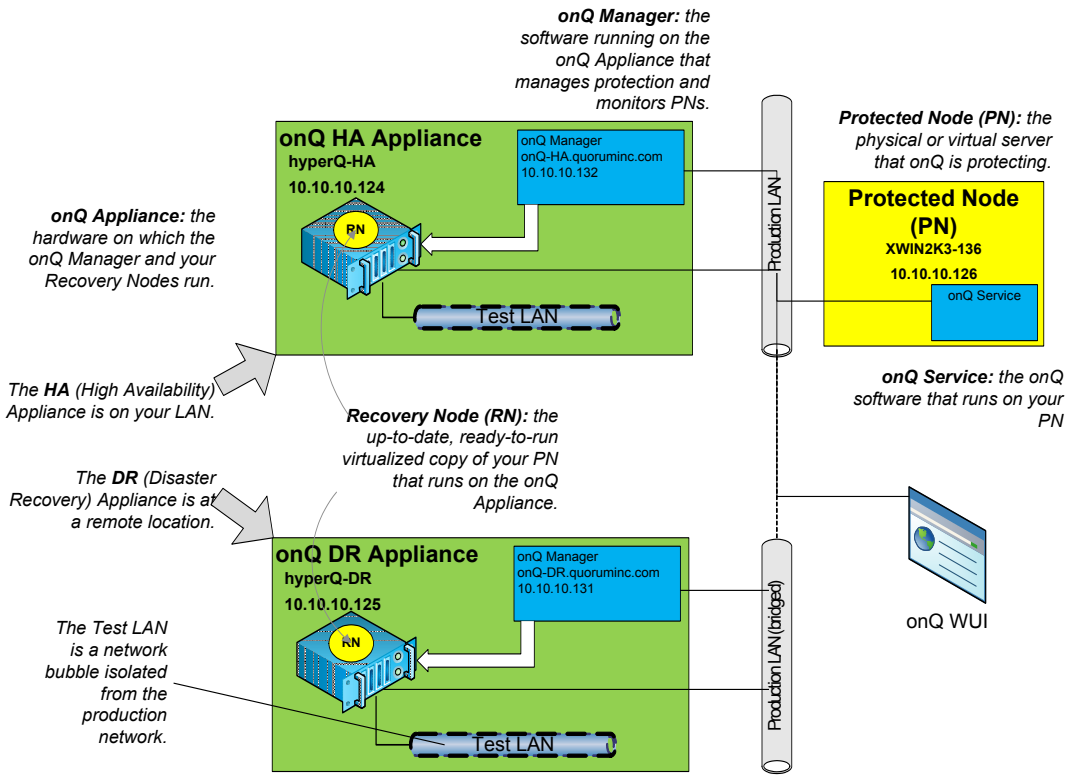
onQ is based on Quorum' feature-rich management and automation platform and comes complete with all the technologies necessary to provide high-availability and disaster-recovery protection for any Windows-based workload. With its intuitive user interface onQ is simple to install and manage.

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:



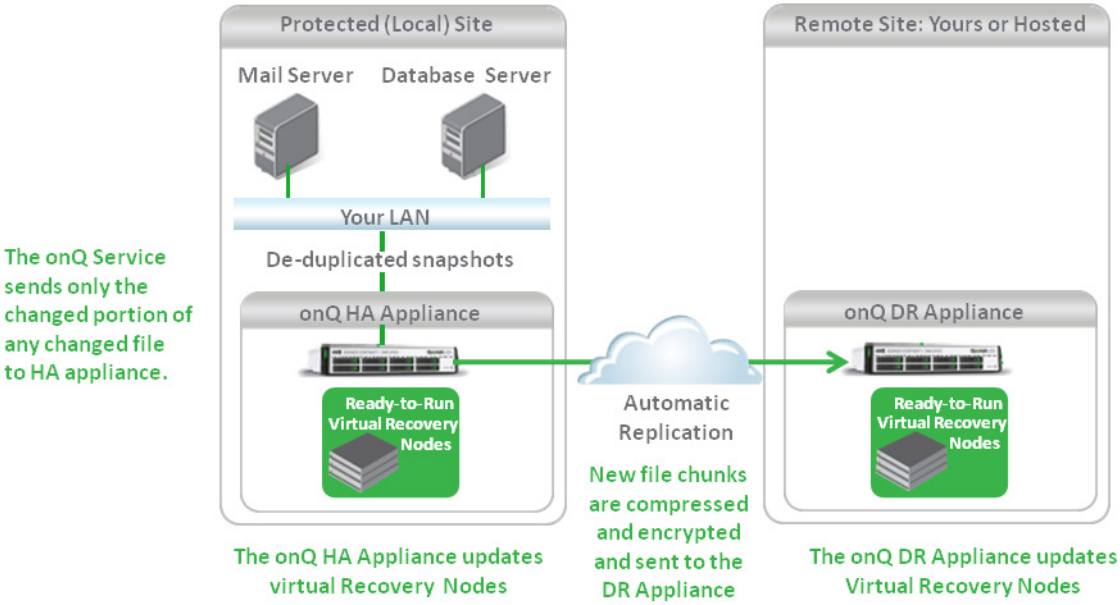
1.3 How does onQ work?

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.
- **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.
- 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

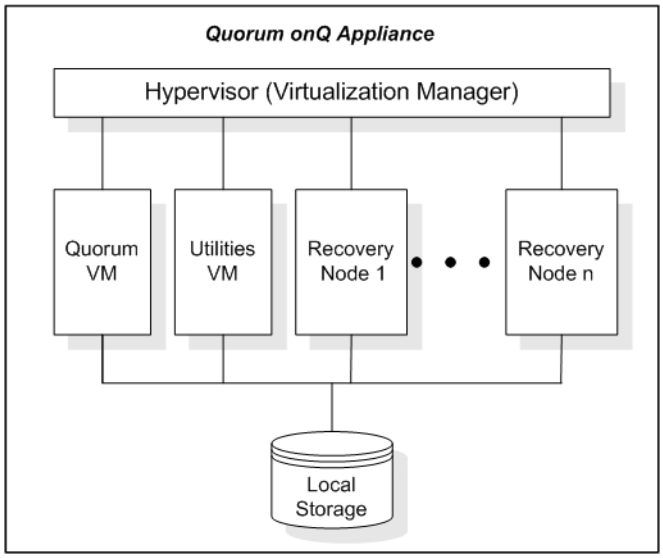
Each Quorum onQ Appliance is a server class x86-based computer running



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 in the cloud. The Deployment of a Hybrid Cloud-enabled HA Appliance is virtually identical to a non-cloud DR Appliance. The only difference is that you must also add information so that the HA Appliance can VPN to the DR Appliance. This secure connection enables the HA Appliance to send snapshots to the DR Appliance.

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. From this onQ Appliance, you can perform the needed BMR/failback to restore your PNs.

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.

1.6 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 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.
- **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.

2.2 Network and Firewall Requirements

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

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. 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 & Outbound	443	TCP	https from onQ IP needed to the Internet for licensing and alerting
Outbound	443	SSL & SMTP	This port is normally SSL/https
Outbound	443	SSH	Communication with License Server, (lic.onqcentral.com)
Outbound	443	SMTP	Sending of alerts (alerts.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:

Port	Protocol	Purpose
3000	TCP	squirtcopy
5000	TCP	wvhds – FLR
5990	TCP	acd control
5991	UDP	PN monitoring
5992	UDP	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:

Port	Protocol	Purpose
22	TCP	may be remapped
81	TCP	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:

Port	Protocol	Purpose
22	TCP	<DR IP> to <HA IP>
81	TCP	<DR IP> to <HA IP>

Port	Protocol	Purpose
80	TCP	Deny ANY PUBLIC / External / Internet to <DR IP> (HTTP protocol)
443	TCP	(https) Allow <DR IP> outbound to Public / External / Internet
443	SMTP & SSH	Allow SMTP, SSH protocols to alerts.onqcentral.com & lic.onqcentral.com respectively

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

Do the following:

- Create local firewall & 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 with 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.

Moreover, in the event of an HA site disaster, the boot order for your recovery nodes must take into account these dependencies.

3

Installation

- 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 Appliances are architecturally identical, although they can have different physical resources available, such as processors and memory. If your installation uses two 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 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 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 Troubleshoot 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. 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. Enroll the Protected Nodes you want to protect.	Step 5: Enroll protected nodes
7. Start protection on the nodes.	Step 6: Start automated protection process
8. (Optional) Install the Archive Vault Storage Module, then configure 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.

3.2.1 Step 1: Prepare to install

The setup process requires Microsoft Internet Explorer 8.0 or greater.

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**.

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

   b. Specify the parameters, including networking information and onQ Role, then **SAVE**:

      The **Modify Remote onQ Setup** page appears.

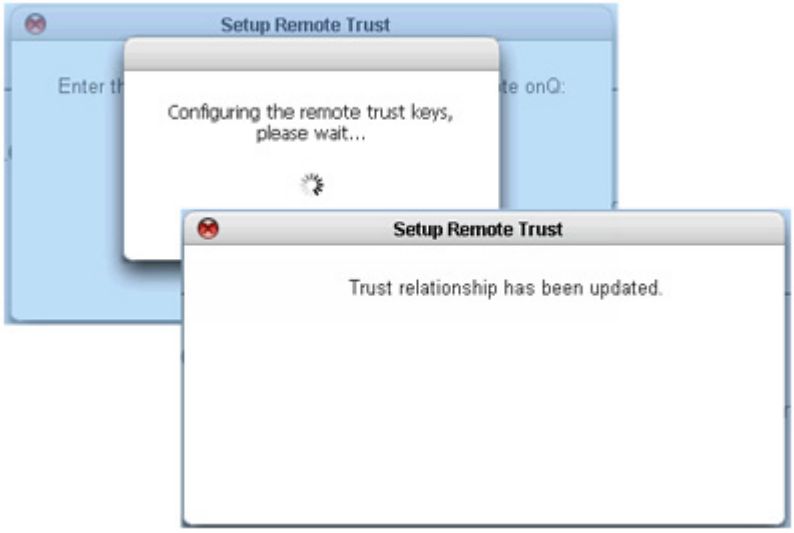
   For examples:
   - Local: Example of DR Appliance
   - Local: Example of DR Mirror
7. Specify remote onQ Appliance configuration:

   Click **MODIFY** and set the parameters, including networking information, DR mirroring, and bandwidth limits, then **SAVE**.

   The **Set Up Remote Trust** page appears.

   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:
   a. Type the VARAdmin password for the *other* onQ Appliance, then **Update Trust**.

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.



b. Test the link.

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

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

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

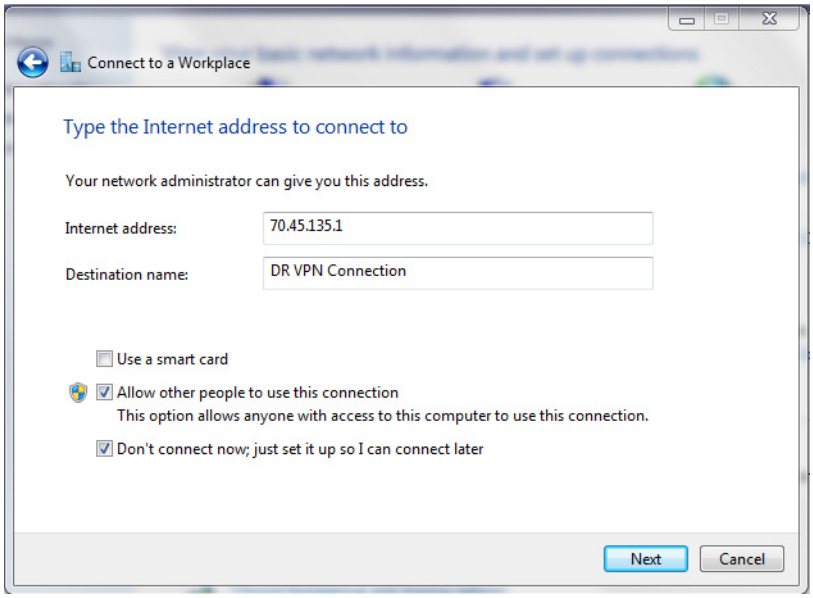
support@quorum.net.

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.

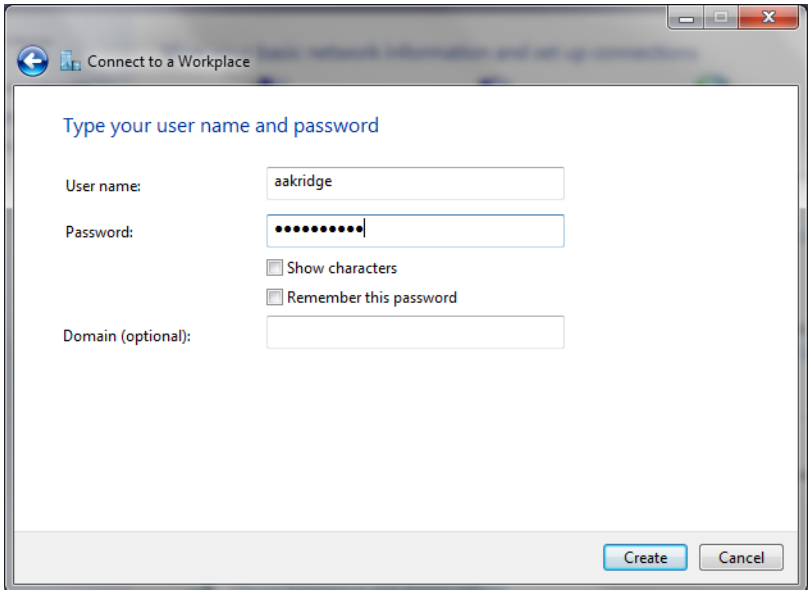
1. Retrieve the following information:
   - Public IP address and logon credentials of the DR Appliance. This information was provided by your Quorum Support Engineer.
   - LAN IP address of DR Appliance in the Quorum Data Center
2. Configure your network firewall to allow PPTP Outbound Traffic.
3. Run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



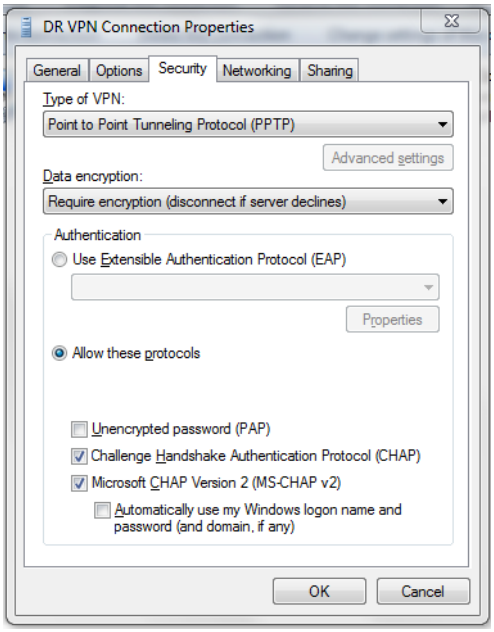
9. Configure the connection to use PPTP protocol:

   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).`

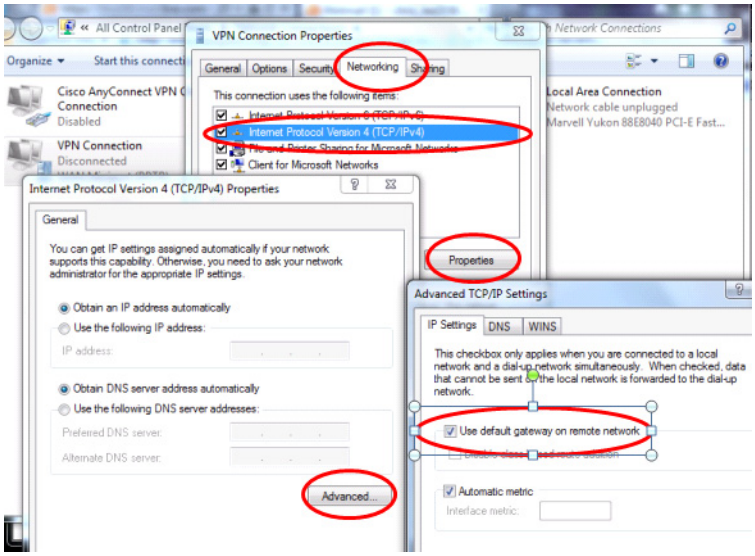
d. Accept all remaining defaults. No further settings are required, then **OK**.



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>`.

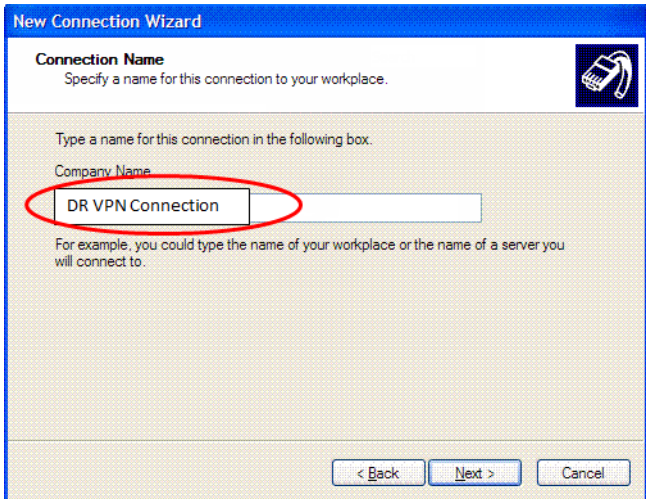
    You should now have full access to the RNs and File Level Backups on the DR Appliance.

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.

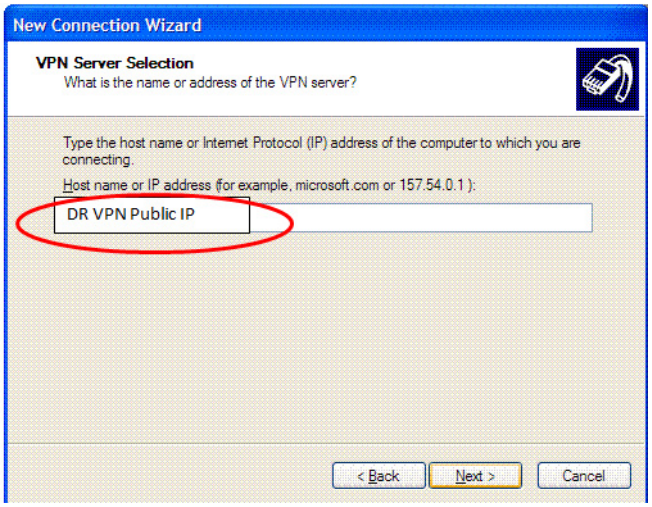
1. Retrieve the following information:
   - Public IP address and logon credentials of the DR Appliance. This information was provided by your Quorum Support Engineer.
   - LAN IP address of DR Appliance in the Quorum Data Center.
2. Configure your network firewall to allow PPTP Outbound Traffic.
3. 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**.



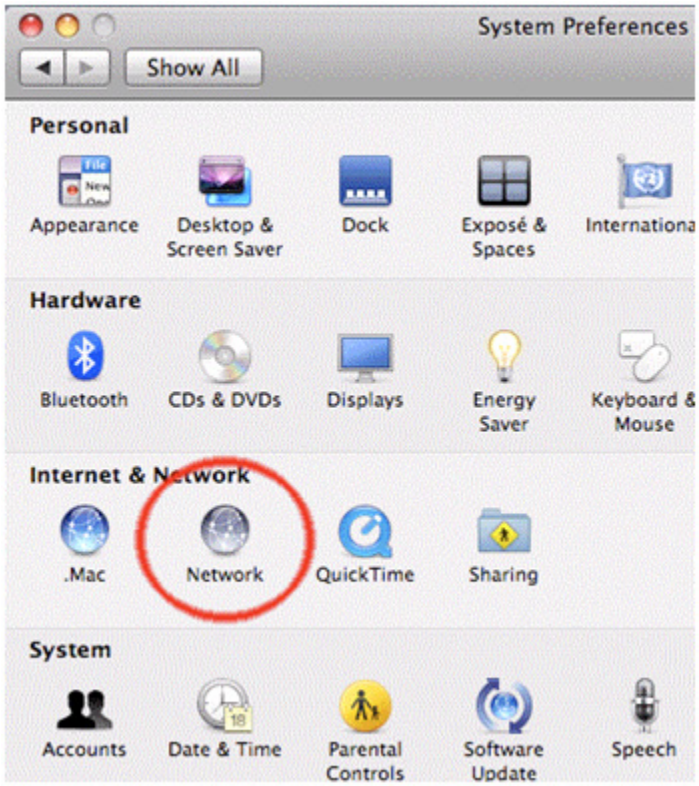
11. If prompted select either **Anyone's use** or **My use only**, then Next, then **Finish**.
12. Modify the connection's settings:
    a. Click on the **Properties** button.
    b. Click on the **Networking** tab.
    c. Highlight **Internet Protocol (TCP/IP)**, then **Properties**.
    d. Click on 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.

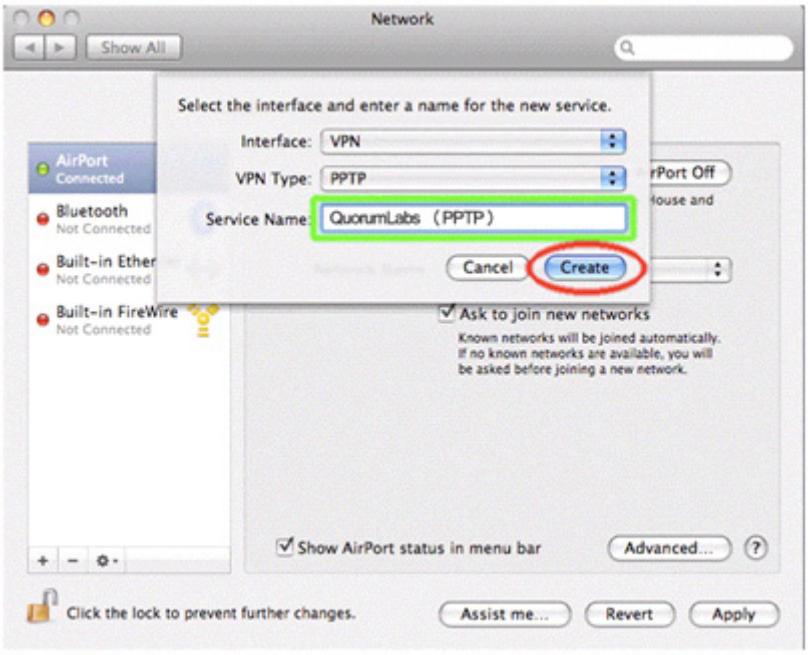
A new service sheet appears.

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**.

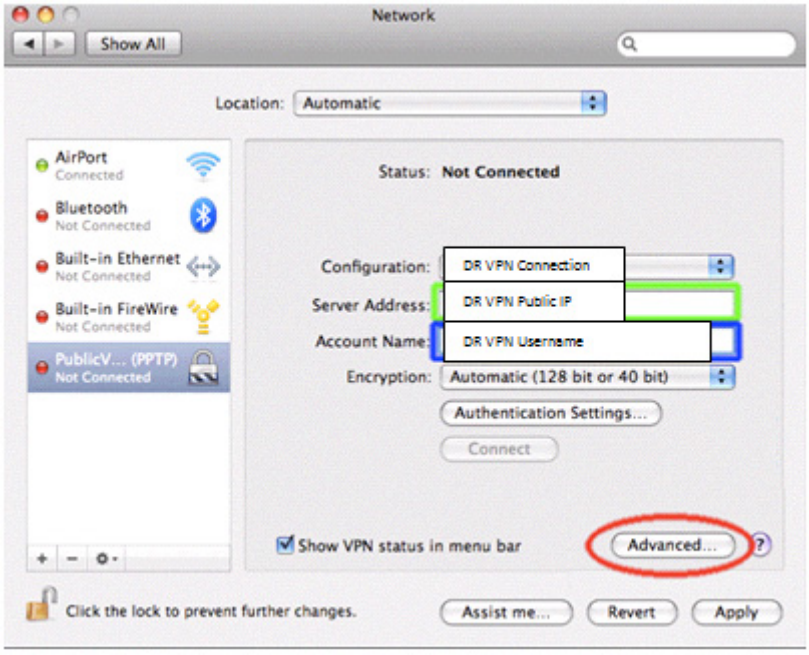
If you set up the VPN previously, the configuration appears in the list.

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

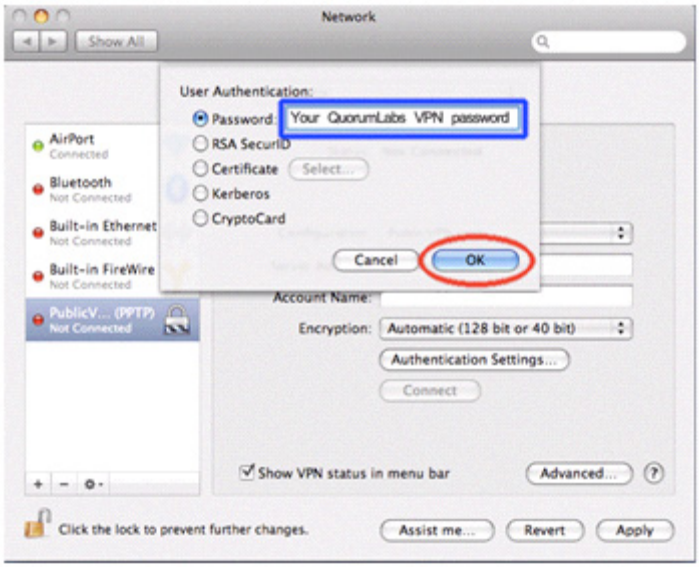
c. In the **Server Address** and **Account Name** fields, specify the Quorum IP address (WAN IP address) and Quorum username.

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.

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: Enroll protected nodes

3.2.5 Step 5: Enroll protected nodes

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

- (Windows) Enroll protected nodes
- (Linux) Enroll protected nodes

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 6: Start automated protection process

3.2.6 Step 6: 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 initial backup begins based on the initial backup time that you scheduled during your setup. See Step 5: Enroll protected nodes for details on setting the initial backup time.

You can monitor progress of the initial backup. See Monitor protected nodes and Monitor DR Appliance.

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, though, 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 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

Updates

- Upgrade Support
- Update Appliance software
- (Windows) Update node software
- (Linux) Update node software
- Get software packages
- Delete software packages
- Download software node packages
- Install onQ Appliance license
- View onQ software versions

4.1 Upgrade Support

Upgrading directly to version 3.6 HF1 is supported for onQ Appliances and protected nodes for 3.6 only. Directly upgrading to 3.6 is supported for the following versions:

- 3.2SP2, GA and HF3
- 3.2SP3
- 3.4, 3.4HF1, 3.4HF2, 3.4HF3, and 3.4HF4

When upgrading, take note of the following:

- After an 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.
- You must manually reboot the onQ Appliance after the upgrade completes.
- Neither 3.6 nor 3.6 HF1 involves any major repository migration or restructuring.

- 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.

4.2 Update Appliance software

Quorum releases a new version of appliance software, service pack and hot fixes as required from time to time. If the onQ appliance is registered with Quorum's onQ Central system for licensing, alert redirection and receiving updates, then updates are automatically downloaded on to the onQ appliance between 10pm and 2am when the updates become available.

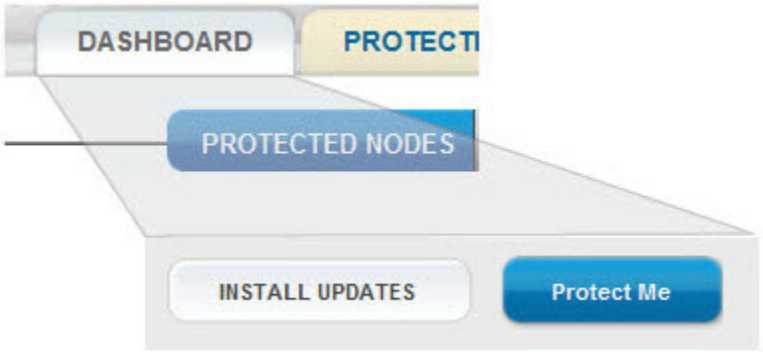
To *automatically* install an update:

Before you begin, verify that your upgrade path is supported.

1. 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 onQ Archive VaultonQ Appliance software downloads from onQ Central.



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. The length of time that onQ requires to complete the upgrade depends on the number of active jobs currently running.

2. 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.

   From the **DASHBOARD**, mouse-over the Install Updates icon to display the status of the upgrade.

3. Restart protection.

4. From your browser, press CTRL-F5 to load and activate the onQ Portal changes.

To *manually* install an update:

You must log on as `VARAdmin` to update the software on your Appliance(s).

- If protection was OFF, then the installation process begins to occur 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.
- If protection was ON, then onQ sends a request to quiesce the system before starting the installation. The installation process waits until onQ is

completely paused before proceeding. Graciously stopping all activities on the onQ takes time. After the installation starts the protection is turned OFF. Protection is restarted after the installation completes and onQ functions resume.

1. Verify that your upgrade path is supported. Go to Upgrade Support.
2. Log on to the onQ Portal.
3. Stop protection.
4. Upload the package.
5. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **Software Upgrades** page.
6. Select the file.
7. Click **Install**.
8. From your browser, press CTRL-F5 to load and activate the onQ Portal changes.

4.3 (Windows) Update node software

You do not need to initiate an upgrade of the onQ Service on your PNs because onQ does this for you automatically 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.

During such an upgrade, the Event Log shows the coordination that transpired between onQ and the protected nodes.

There is no way to disable automatic upgrades of the onQ Service. However, you can continue to use the following two alternative upgrade methods.

To *automatically* install an update:

You must log on as `VARAdmin` to update the onQ Service on your PNs.

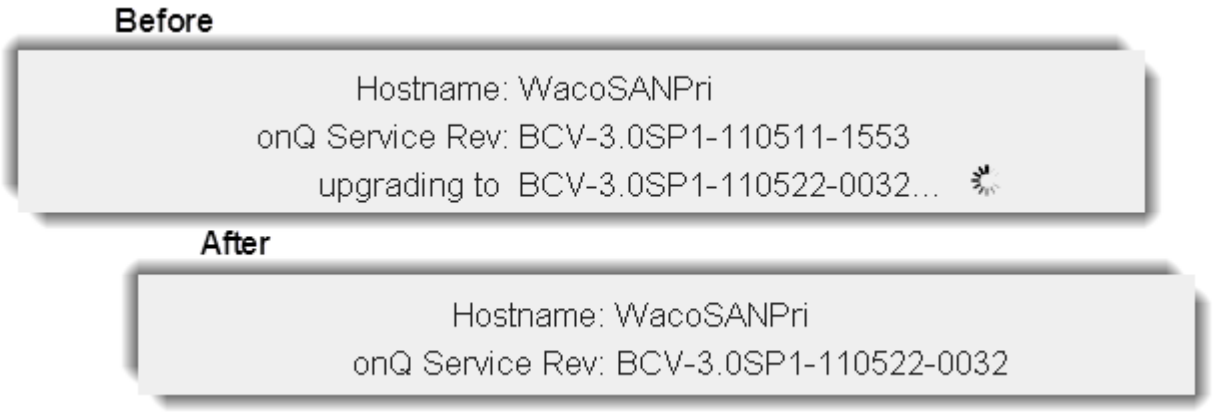
1. Verify that your upgrade path is supported. Go to Upgrade Support.

2. Log on to the onQ Portal.
3. Download the package.
4. Ensure that protection is On.
5. Go to **PROTECTION CONFIG** tab.
6. Select the PN that you want to upgrade, then **VIEW** button.
7. From the View a Protected Node dialog, click **Upgrade onQ Service** button. This button does not appear if your Protected Node is running the latest software version.

   A `Submitting Upgrade Request` message appears along with a progress icon.

   As the PN is upgrading, the View a Protected Node dialog shows the current version of the onQ Service running on the PN, and the version it will be upgraded to.

   After the upgrade completes, the new version appears in the **onQ Service Rev** field:



To *manually* install an update:

You must log on as `VARAdmin` to update the onQ Service on your PNs.

1. Verify that your upgrade path is supported. Go to Upgrade Support.
2. Log on to the onQ Portal using Internet Explorer. If you want to use Firefox, download the software before you continue.
3. Upload the package.
4. Stop protection.

5. Click **Protect Me**.The Protect Me wizard launches.

   The wizard automatically finds the package that you uploaded.

6. Follow the wizard's instructions to upgrade the PNs.

Related Topics

Upgrade Support
(Linux) Update node software

4.4 (Linux) Update node software

This procedure requires that you log on as either onQ varadmin or admin user. Only these users can install/update the node packages on Linux PNs. During the install the Linux PN needs to communicate with onQ to send/retrieve information to enroll/update PNs.

To update node software:

1. Log on to the Linux PN.
2. (64-bit PNs running 3.4, 3.4HF1, and 3.4HF3). Due to the known limitation outlined in "Linux RN Limitations" in onQ Release Notes, find and remove the existing node packages.

```
# pkgs=`rpm -qa|grep BCV`
    # if [ "$pkgs" != "" ]
        then
            rpm -e $pkgs
   finish
```

3. (Redhat 6.x PNs with EFI boot) Do one of the following to ensure that the PN can boot. `/boot/grub/grub.conf.xvf5` has higher priority than `/boot/grub/menu.lst`.
   - Copy and modify `/boot/grub/menu.lst`:

```
[root@RHEL62x64-16-110 grub]# ls
/boot/efi/EFI/redhat/grub.conf \
/boot/efi/EFI/redhat/grub.conf


cp /boot/efi/EFI/redhat/grub.conf
/boot/grub/
cd /boot/grub/
ln -s ./grub.conf menu.lst
```

   - Create `/boot/grub/grub.conf.xvf5`.
4. (Redhat 5.x PNs) If this is the first time you're enrolling this PN, install kernel-xen RPM package.

5. From within a folder where you want to store and use the install script, do the following:

   a. Delete any existing `install.py` files.

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

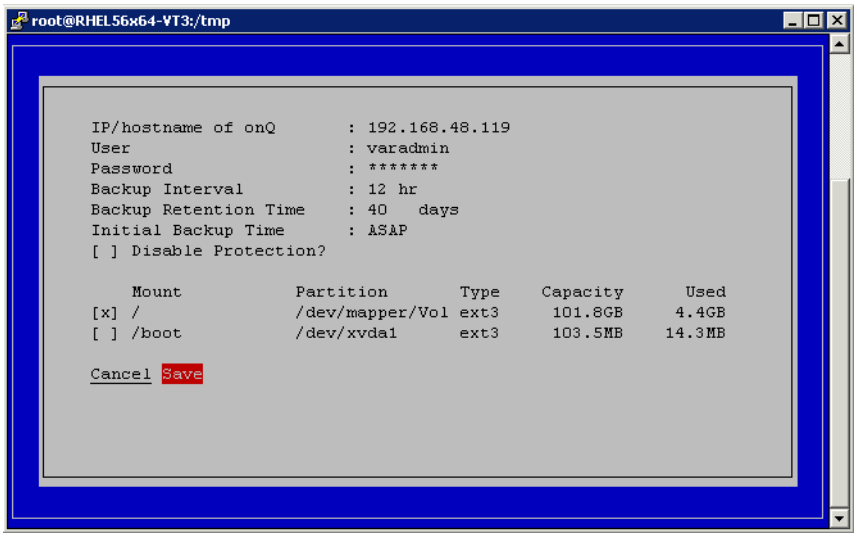
   b. Run the following command to retrieve the installer:

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

6. Run the following command to start the installer:

```
# python ./install.py
```

7. From the utility, type the credentials for either onQ varadmin or admin user.



8. Specify all valid values in the blank fields provided and modify defaults as needed. The install utility lists all available volumes/partitions/mount points.
   - For any Linux PN, you must protect both `/` and `/boot` and any other mount points that contain the necessary software to boot the RN.
   - For PNs running RHEL 5.x and due to the known limitation outlined in "Linux RN Limitations" in onQ Release Notes, set the

Format (file system type) parameter for both the root partition (/) to ext3.

- For more information on the parameters, go to Protected Node Parameters.

9. Do the following:
   a. Select **Save** to update/install the client node package on the Linux node.
   b. Wait! Do not press Enter. If the packages installed successfully, exit the script: tab down and click **Cancel**.

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 as outlined in Step 10.
Incorrect/invalid values entered	Install utility stops if you type incorrect/invalid values. Correct the problem and **Save** again or **Cancel**.
PN configuration failed: Operation not permitted while protection is running	Turn off protection (go to Stop node protection). Press Enter for the GUI's curser to reappear. Utility preserves all values entered previously.

10. On the protected node, open the following ports:
    - UDP port 5990
    - TCP ports 5000 and 5990

    These ports must be open in order for onQ to see the protected node.
11. (Oracle database) Install RMAN scripts.
12. 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.

Related Topics

(Windows) Update node software

4.5 Get software packages

Before you can install onQ Appliance or PN software updates, you must do one of the following to get the updates:

- 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. Contact Quorum Support for assistance.

After you upload the new software, consider deleting the old version.

To automatically download packages:

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 install them as outlined in Update Appliance software.

1. Log on to the onQ Portal.
2. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **SOFTWARE UPGRADES** page > **Check for Updates** button.

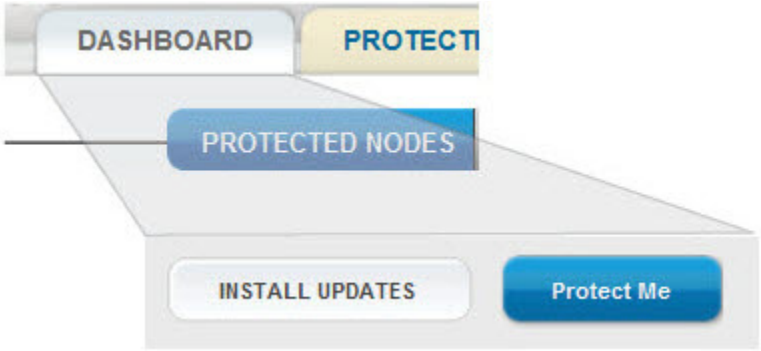
   onQ contacts onQ Central and downloads the latest updates. If a direct upgrade path is not supported for your version, onQ automatically downloads all the prerequisite updates so as to migrate your Appliance and protected nodes to the latest version.
3. Mouse-over the **Check for Updates** button to 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:

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 ApplianceAV 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

4.6 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 (**-**).

4.7 Download software node packages

Normally, there's no need to download the node package because the Protect Me wizard automatically retrieves the software update for you, as outlined in To manually install an update:).

However, if you manually update a PN by using Firefox instead of IE, you need to select the node package that matches the operating system because Firefox will not automatically select the correct 32-bit or 64-bit operating

system version.

To download a node package:

1. RDP to the protected node.
2. From that PN, launch the Appliance's onQ Portal.
3. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **Downloads** page.
4. Select the file.
5. Click **Download**.

4.8 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 Alerts

4.9 View onQ software versions

As outlined in How does onQ work?, onQ is comprised of several components. onQ is aware of all the versions for each component that a given onQ Appliance is running, although onQ doesn't disclose the version of the overall installation package that the onQ Appliance is running (for example, 3.6 HF1). Only Quorum Support can determine the package that the onQ Appliance is running.

However, 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.

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 as onQ performs a software inventory.

5

The Basics

- Log on to Appliance's user interface
- Establish a VPN connection to Hybrid Cloud
- Log off from Appliance's user interface
- Pull-down menu
- Reboot onQ Manager
- Reboot Appliance
- Shut down Appliance
- Modify Appliance's network settings
- Configure Appliance's hypervisor settings
- Unlock UI
- Add protected nodes manually
- Modify protected nodes
- Delete protected nodes
- Protected Node Parameters
- (Windows) Enroll protected nodes
- (Linux) Enroll protected nodes

5.1 Log on to Appliance's user interface

You need a web browser to access the onQ Appliance's user interface. Firefox is recommended. Internet Explorer 7.0 or above is also supported. However, installations and upgrades work best with Internet Explorer.

- 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.

5.2 Establish a VPN connection to Hybrid Cloud

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.

5.3 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 pull-down menu.
2. Choose **Log Off**.

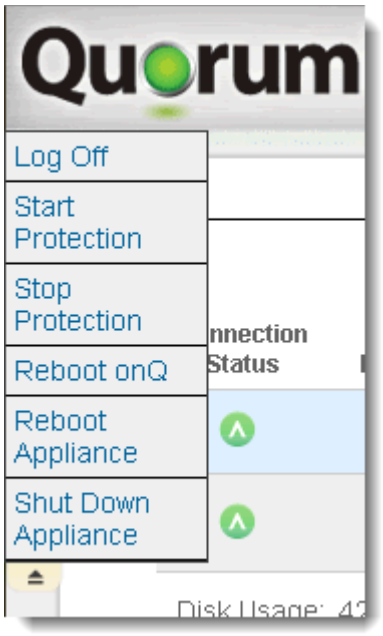
5.4 Pull-down menu

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

5.5 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.

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 pull-down menu.
3. Choose **Reboot onQ**.
4. Click **Continue** to initiate the reboot.
5. After the onQ Manager reboots, verify that protection is on.

5.6 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 pull-down menu.
3. Choose **Reboot Appliance**.
4. Click **Continue** to initiate the reboot.

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

5.7 Shut down 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 Appliance:

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

To restart the Appliance:

1. Ensure that the power is off, then turn it on (power-cycle).
2. Restart protection.
3. Restart any recovery nodes that are in use.

5.8 Modify Appliance's network settings

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

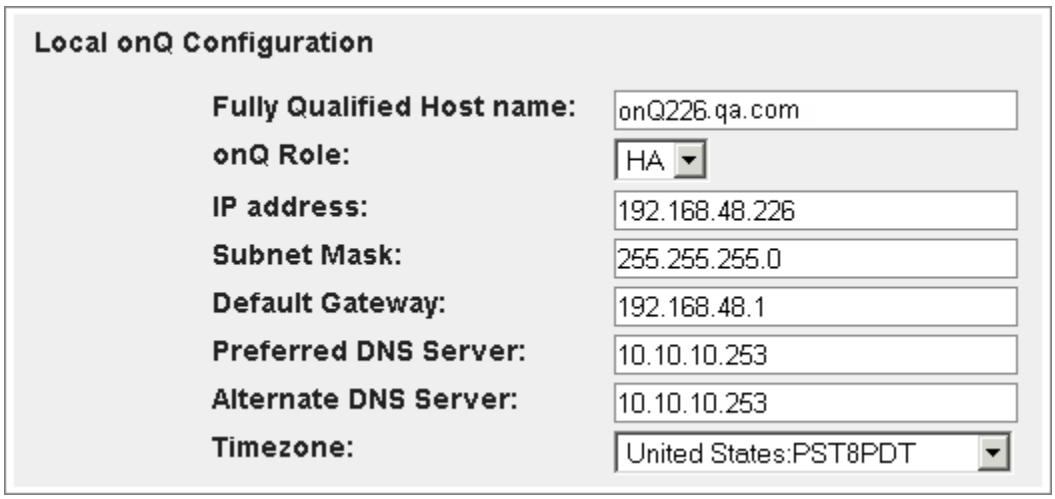
- Fully qualified host name
- IP Address
- subnet mask
- default gateway
- Preferred and Alternate DNS servers.
- onQ Appliance's role (see Change Appliance's role)
- local time zone

To modify an onQ Appliance's network settings:

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 **RESET** to reset 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.

5.9 Configure Appliance's hypervisor settings

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

- name
- IP address
- password
- subnet mask
- default gateway
- preferred DNS server address
- alternate DNS server address
- NTP server

- time zone

To configure an onQ Appliance's hypervisor settings:

1. Log on to the onQ Appliance's onQ Portal.
2. Click the **APPLIANCE CONFIG** tab > **HYPERVISOR** page > **MODIFY** button.

   The **Modify Hypervisor Config** page appears.
3. Make your changes, then click **SAVE**.

   Click **RESET** to reset 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.

5.10 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:

5.11 Add protected nodes manually

When you enroll protected nodes, onQ automatically adds them to the onQ Manager, among other things; therefore, you don't need to add them manually. However, there are reasons why you might want to add a protected node manually:

- 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 nodes yet.

To add a 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. Specify temporary values for the node parameters.
6. Click **Yes** to confirm the request.

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



7. In the **Backup Exclude List** dialog, click **DEFAULT**. You can add backup exclusions at any time. If the **File Backup Only** option was selected, then in the **Backup Include List** dialog, click **Custom**. You can add or modify backup inclusions at any time
8. From the **PROTECTION CONFIG** tab, confirm that the newly added protected node appears in the list.
9. If you'd eventually like to protect this node, you must stop protection, then enroll the protected node.
10. Restart protection.

5.12 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.

The **PROTECTION CONFIG** tab on a DR Appliance is slightly different from its HA counterpart. The DR version shows the HA Appliance's backup retention time as well as its own backup retention time. The DR version does not show a protection disabled status because this concept does not exist on a DR Appliance.

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.

5.13 Delete protected nodes

During the setup process you identified the nodes you want to protect. As you remove hosts from production, you might want to remove a node 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.

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

Note: Removing a protected node deletes the protected node’s virtual machine (vdisk), not its snapshots. To delete snapshots, go to Delete snapshots.

To remove a node:

1. Log on to the HA Appliance’s onQ Portal.
2. Stop protection.
3. From the list of protected nodes, select the node from the list.
4. Click the button labeled with the minus sign (**-**).
5. Click **Yes** to confirm the request and return to the **PROTECTION CONFIG** tab.

   The Protected Node disappears from the list.
6. Restart protection so that onQ recognizes that you’ve deleted the node.

5.14 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, enter the host name or IP address of the node to protect. This field is display-only under **MODIFY**.
- **Backup Interval**: The interval (in minutes/hours) between the start of data backups. To monitor these backups, go to "Monitor protected nodes" on page 175 and "Monitor backups" on page 181.
- **Backup Retention Time**: The maximum number of days to retain backed-up data that is no longer referenced. The **PROTECTION CONFIG** tab on a DR Appliance is slightly different from its HA counterpart: the DR version shows the HA Appliance's backup retention time as well as its own backup retention time; the DR version does not show a protection disabled status because this concept does not exist on a DR Appliance.
- **Initial Backup Time**: Select the time of day for the initial backup. Select **ASAP** to begin initial backup as soon as possible. Be aware that 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.
- **OS Type?**: Both Linux and Windows operating systems are supported, though the parameters available vary.
- **Auto RN Creation?**: Instructs onQ to create a recovery node for this protected node.
- **Disable Protection?**: Select *yes* or *no* to disable or enable protection for this node. If you have disabled 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 82.

Windows-Only Parameters

There are some parameters of the recovery node that correspond to the protected node:

- **Number of CPUs**: # of CPUs for this recovery node.
- **RN Memory**: RN's memory in GBs.
- **Disk Letter**: Disk's letter or Mount Point.
- **Size**: Size of the disk in GBs.

Linux-Only Parameters

- **Number of CPUs**: # of CPUs for this recovery node.
- **RN Memory**: RN's memory in GBs.
- **Mount Point**: Where the protected node's filesystem is mounted.
- **Format**: File system type.
- **Size**: Size of the disk in GBs.

Advanced Parameters

There are a few **ADVANCED** parameters:

- **Execute Pre Snapshot script?** executes before a snapshot starts on the protected node. See "Run custom backup scripts" on page 121.
- **Execute Post Snapshot script?** Executes soon after a snapshot completes on the protected node. See "Run custom backup scripts" on page 121.
- **Enable Filter Driver?** (This parameter is not applicable for Linux-based protected nodes.) See "Methods of Performing Incremental Backups" on page 126.
- **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?** onQ's backup utility clears VSS flags after it performs backups. This process can conflict with some applications 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 option to **No** (default).

- **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 Windows Cluster Services.) If you select **Yes**, specifies that the node is part of an MSFT cluster (Microsoft Cluster Server/MSCS). MSFT for Windows 2012, Windows 2008 R2, Windows 2008, and Windows 2003 are supported. Only iSCSI connections to a storage device in a SAN environment is supported. 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.
- **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 101.

Related Topics

Modify protected nodes
Add protected nodes manually
(Windows) Enroll protected nodes

5.15 (Windows) Enroll protected nodes

Each 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.

You must perform this procedure from the node that you want to enroll. Repeat this procedure for each node that you want to enroll.

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 the volume configured with REFS that is new in Windows Server 2012.

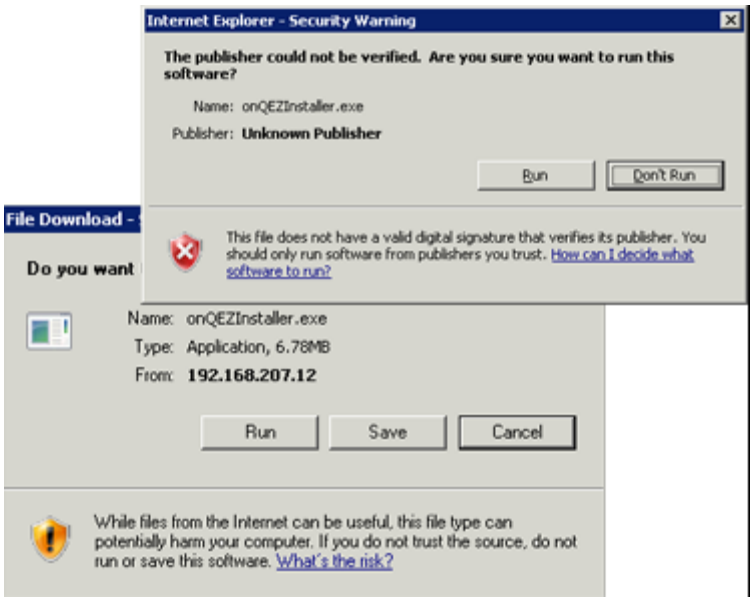
Note: Even if you added the new node manually and it appears in the list of protected nodes, you must enroll the node before the onQ Manager can protect it. Previously protected nodes that are still running the onQ Service do not need to be enrolled; in this case, simply add the nodes manually, then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

To enroll the protected nodes:

1. Log on to the server that you want to enroll as a PN.
2. From that server, log on, as a user with `Administrator` privileges, to the HA Appliance's onQ Portal using a supported browser as outlined in "Browser Support" in onQ Release Notes.
3. Click on the **PROTECTION CONFIG** tab.
4. At the bottom of the page, click the **Protect Me** button, then **CONTINUE**.

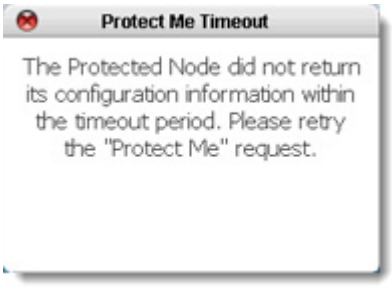
5. Run the onQ Service executable (`onQEZInstaller.exe`).

   a. 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.



   After the installer discovers and displays the disk and memory parameters for the RN corresponding to this PN, the **Modify Protection Parameters** dialog appears behind the **Setup Wizard** launches.

   - 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**.

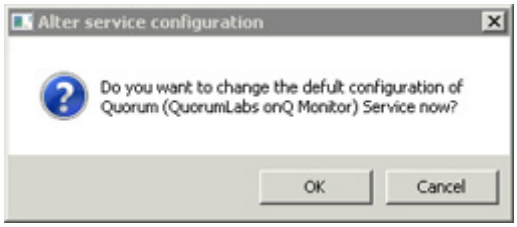


6. Install onQ Service.

   Follow the on-sreen instructions in the **Setup Wizard**.

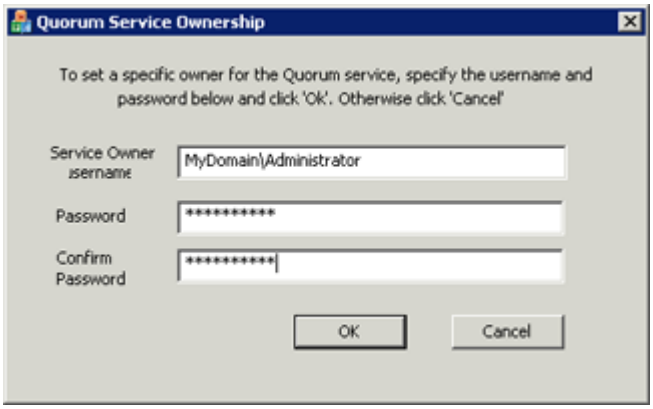


When prompted, choose whether or not you want to configure the service to use a specific account. This account must have adequate privileges to back up all the files that you want to protect.



In the dialog box, 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 RN's proposed configuration.

7. Modify PN parameters.

In the **Modify Protection Parameters** dialog, specify values for the node parameters, then click **SAVE** to add the node to the list of protected nodes.

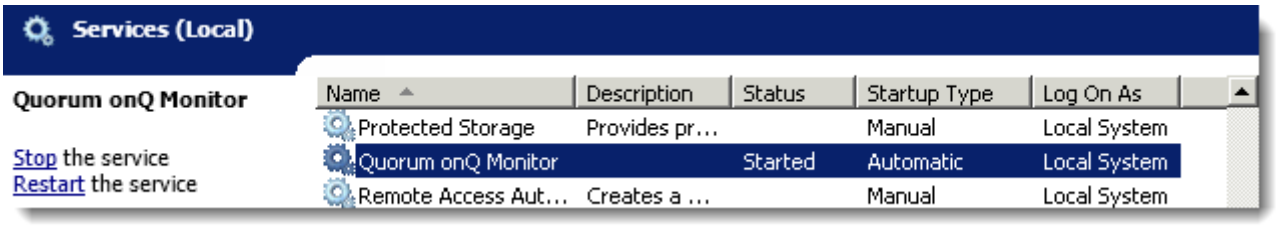
Note: If you don't see the Modify Protection Parameters dialog, move the wizard because it might be hiding the dialog.

Optionally, you can modify the node at any time.

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

8. Verify PN enrollment.

   Verify that the onQ Service is running on the PN:



5.16 (Linux) Enroll protected nodes

This procedure requires that you log on as either onQ varadmin or admin user. Only these users can install/update the node packages on Linux PNs. During the install the Linux PN needs to communicate with onQ to send/retrieve information to enroll/update PNs.

To update node software:

1. Log on to the Linux PN.

2. (64-bit PNs running 3.4, 3.4HF1, and 3.4HF3). Due to the known limitation outlined in "Linux RN Limitations" in onQ Release Notes, find and remove the existing node packages.

```
# pkgs=`rpm -qa|grep BCV`
    # if [ "$pkgs" != "" ]
        then
            rpm -e $pkgs
   finish
```

3. (Redhat 6.x PNs with EFI boot) Do one of the following to ensure that the PN can boot. `/boot/grub/grub.conf.xvf5` has higher priority than `/boot/grub/menu.lst`.
   - Copy and modify `/boot/grub/menu.lst`:

```
[root@RHEL62x64-16-110 grub]# ls
/boot/efi/EFI/redhat/grub.conf \
/boot/efi/EFI/redhat/grub.conf


cp /boot/efi/EFI/redhat/grub.conf
/boot/grub/
cd /boot/grub/
ln -s ./grub.conf menu.lst
```

   - Create `/boot/grub/grub.conf.xvf5`.
4. (Redhat 5.x PNs) If this is the first time you're enrolling this PN, install kernel-xen RPM package.
5. From within a folder where you want to store and use the install script, do the following:
   a. Delete any existing `install.py` files.

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

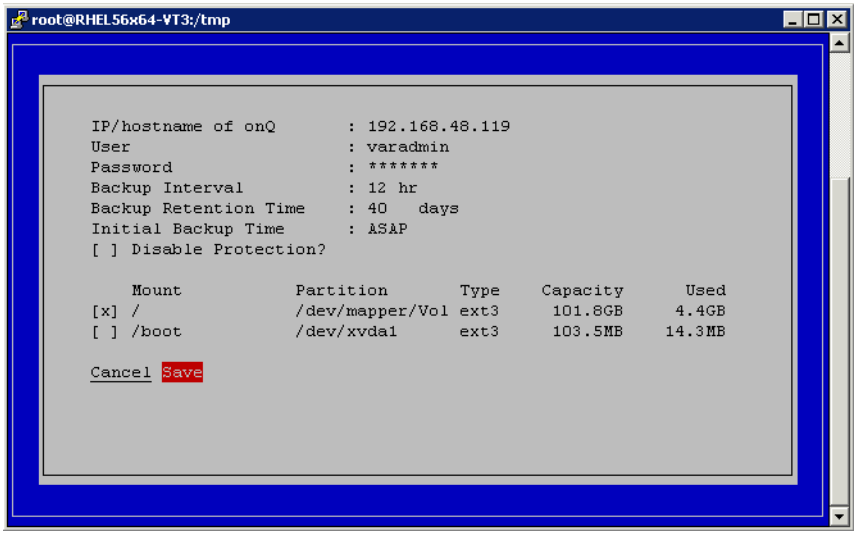
b. Run the following command to retrieve the installer:

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

6. Run the following command to start the installer:

```
# python ./install.py
```

7. From the utility, type the credentials for either onQ varadmin or admin user.



8. Specify all valid values in the blank fields provided and modify defaults as needed. The install utility lists all available volumes/partitions/mount points.
   - For any Linux PN, you must protect both `/` and `/boot` and any other mount points that contain the necessary software to boot the RN.
   - For PNs running RHEL 5.x and due to the known limitation outlined in "Linux RN Limitations" in onQ Release Notes, set the **`Format`** (file system type) parameter for both the root partition (`/`) to `ext3`.
   - For more information on the parameters, go to Protected Node Parameters.

9. Do the following:

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

b. Wait! Do not press Enter. If the packages installed `successfully`, exit the script: tab down and click **Cancel**.

`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 as outlined in Step 10.

`Incorrect/invalid values entered`	Install utility stops if you type incorrect/invalid values. Correct the problem and **Save** again or **Cancel**.
`PN configuration failed: Operation not permitted while protection is running`	Turn off protection (go to Stop node protection). Press Enter for the GUI's curser to reappear. Utility preserves all values entered previously.

10. On the protected node, open the following ports:
    - UDP port `5990`
    - TCP ports `5000` and `5990`

    These ports must be open in order for onQ to see the protected node.
11. (Oracle database) Install `RMAN` scripts.
12. 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.

6

Protection

- Start node protection
- Stop node protection
- Modify recovery node creation policy
- Recreate recovery nodes
- Disable protection of nodes
- Create on-demand RNs
- Delete on-demand RNs
- Cancel RN updates
- Manage hosts
- Resize protected node’s vdisk
- Disable snapshot transfers
- Enable and disable DR Mirroring
- Enroll and Disenroll onQ Archive Vault

6.1 Start node protection

After you add nodes, then enroll them, you're ready to protect them.

It will take several hours to complete the initial backup cycle and recovery node construction, 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 pull-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:



6.2 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.

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 pull-down menu.

4. Choose **Stop Protection**, then **Stop**.

   The **Protection:OFF** icon turns red:



6.3 Modify recovery node creation policy

onQ can create recovery nodes either automatically with each backup cycle or on-demand when requested.

Recovery nodes can be configured to be built in two ways:

- **Ready-to-Run** - the RN is updated after each backup cycle
- **On Demand** - the RN is only built when requested through the onQ Portal

Build-on-Demand mode is typically used to conserve onQ Appliance resources by creating RNs less frequently and only as they are needed.

To switch policies: (RECOVERY NODES page)

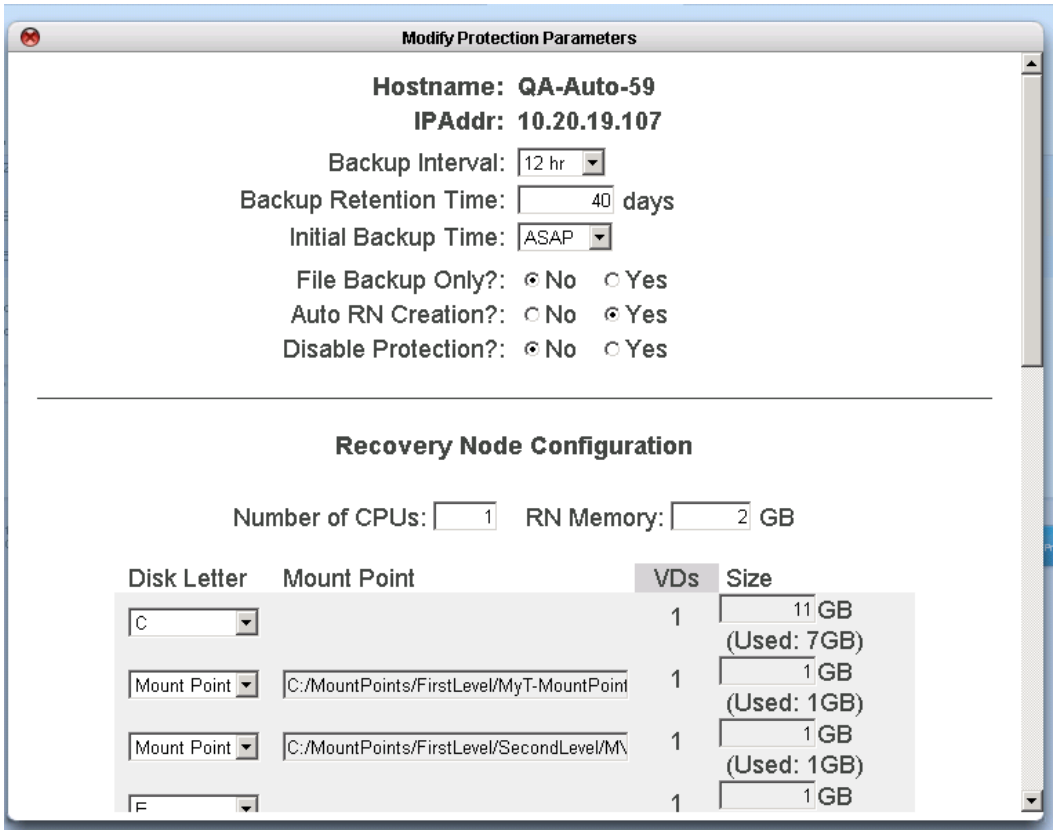
1. Log on to the HA Appliance's onQ Portal.
2. Go to **DASHBOARD** tab > **RECOVERY NODES** page.
3. Unlock the page.
4. Do one of the following:
   - (All RNs) Click either the **ALL BoD** or **ALL AUTO** button at the bottom of the page, then click **Disable RN Auto Creation/Enable RN Auto Creation** 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 icon if `Auto RN Creation` is set to disabled.

To switch policies: (PROTECTION CONFIG tab)

1. Log on to the HA Appliance's onQ Portal.
2. From the **PROTECTION CONIG** tab, select the PN you want to modify, then click **MODIFY**.
3. Change the Auto RN Creation setting, then click **SAVE**.
   - Ready-to-Run - Auto RN Creation=Yes
   - Build on Demand (BoD) - Auto RN Creation=No



When RN Auto Creation is set to No, the RN must be built manually. This can take anywhere from a few minutes to several hours, depending on size. There must be enough disk space on the onQ Appliance to accommodate the RN.

6.4 Recreate recovery nodes

Recreating a recovery node simply means creating a recovery node from a previous snapshot. 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.

To recreate a recovery node (HA Appliance):

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.

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.

6.5 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.

When you re-enable protection, backups resume and subsequent RN rebuilds are based on the retained RN.

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 either the **DISABLE ALL** or **ENABLE ALL** button at the bottom of the page, then click **Disable Protection/Enable Protection** 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.

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**.

6.6 Create on-demand RNs

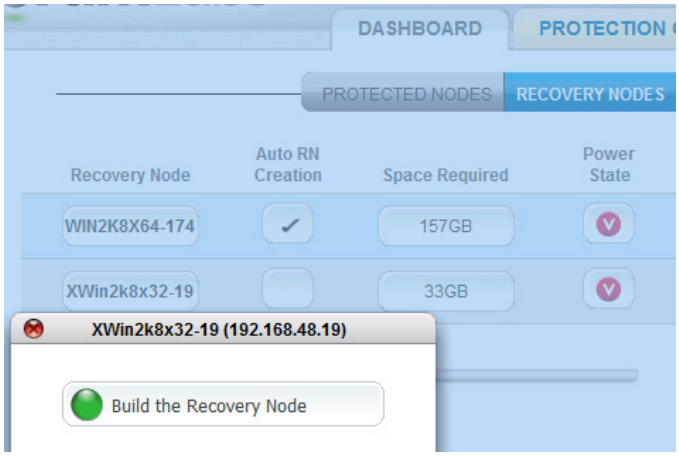
If you specified an On Demand policy, you must create recovery nodes manually. This process can take 20 minutes to several hours depending on the size of the recovery node. onQ uses the recent backup to create the recovery node.

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.



6.7 Delete on-demand RNs

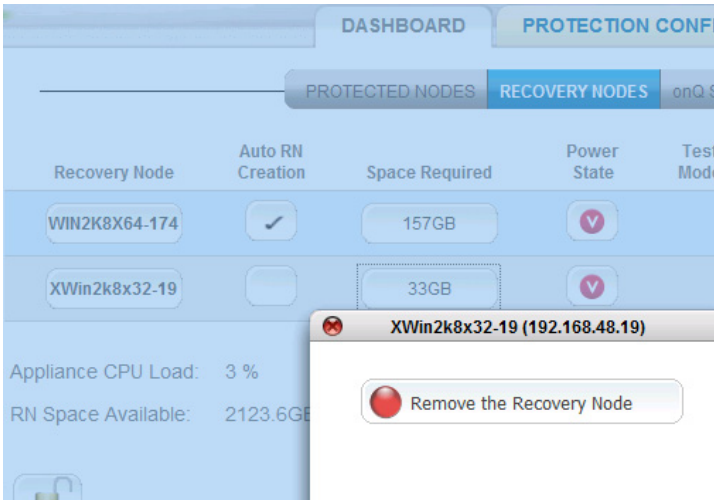
Each RN uses nearly the same amount of disk space on the onQ Appliance as is used on the PN.

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**.



6.8 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 (cycle) 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.



6.9 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 e-mail 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**.

6.10 Resize protected node's vdisk

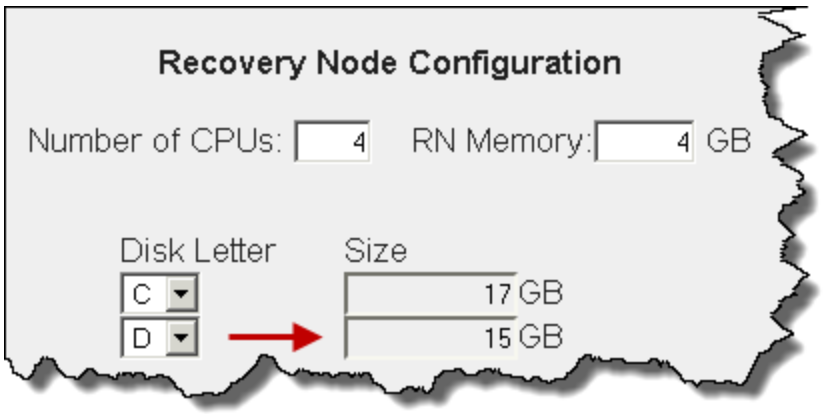
As outlined in "Add protected nodes manually" on page 62, a protected node's virtual disk (vdisk) requires a specific amount of disk space.

As soon as your protected node's vdisk capacity used exceeds 80%, onQ alerts you (see "vDisk or Disk Space Alerts" on page 196). If you wait until a vdisk fills up, onQ cannot protect that protected node.

At any time, you can resize the vdisk without deleting and rebuilding the protected node.

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**:



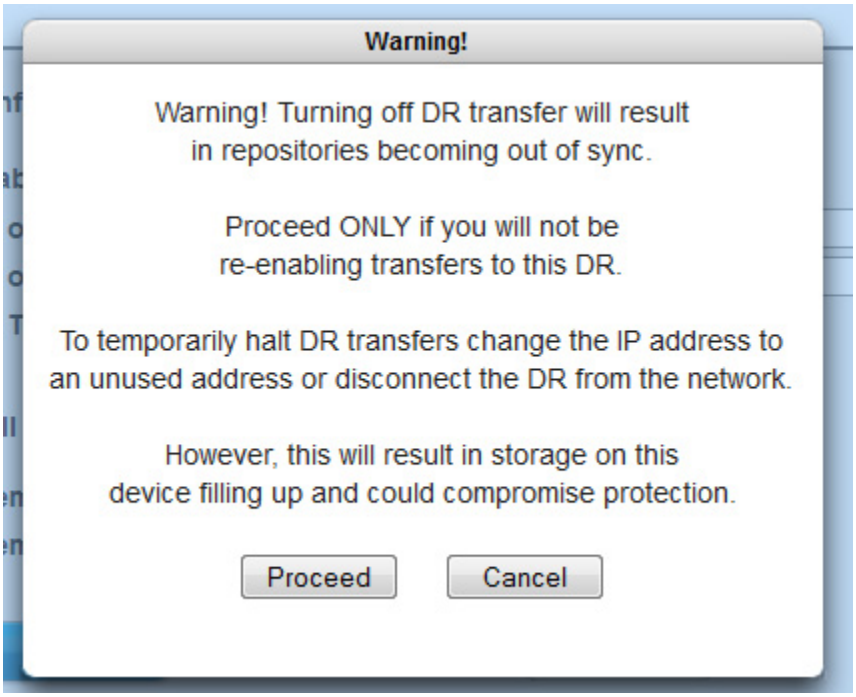
Related Topics

vDisk or Disk Space Alerts
Monitor disk space usage

6.11 Disable snapshot transfers

Disabling snapshot transfers is not recommended, unless your configuration is an HA-only configuration, meaning you don't have a DR Appliance; if you have an HA-only configuration, disabling 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. Before you do so, you must have enough disk space to accumulate backups and change the DR Appliance's IP address to an unused address or disconnect the DR Appliance from the network, as the warning below advises:.



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.

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.

To disable snapshot transfers:

1. Ensure that your onQ Appliance has enough disk space to accumulate backups. For the related alert, go to vDisk or Disk Space Alerts.
2. Change the DR Appliance/DR Mirror's IP address to an unused address or disconnect the DR Appliance/DR Mirror from the network.
3. Log on to the onQ Appliance's onQ Portal.
4. Stop protection.
5. Go to **APPLIANCE CONFIG** tab > **onQ (REMOTE)** page > **MODIFY** button.
6. In the **Disable DR Transfer?** field, select the **Yes** radio button, then **SAVE**.
7. Restart protection.

   Disable Transfer: OFF displays at the top of each UI page on the onQ Appliance.

6.12 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**.

Related Topics

Step 2: Configure the onQ Appliances

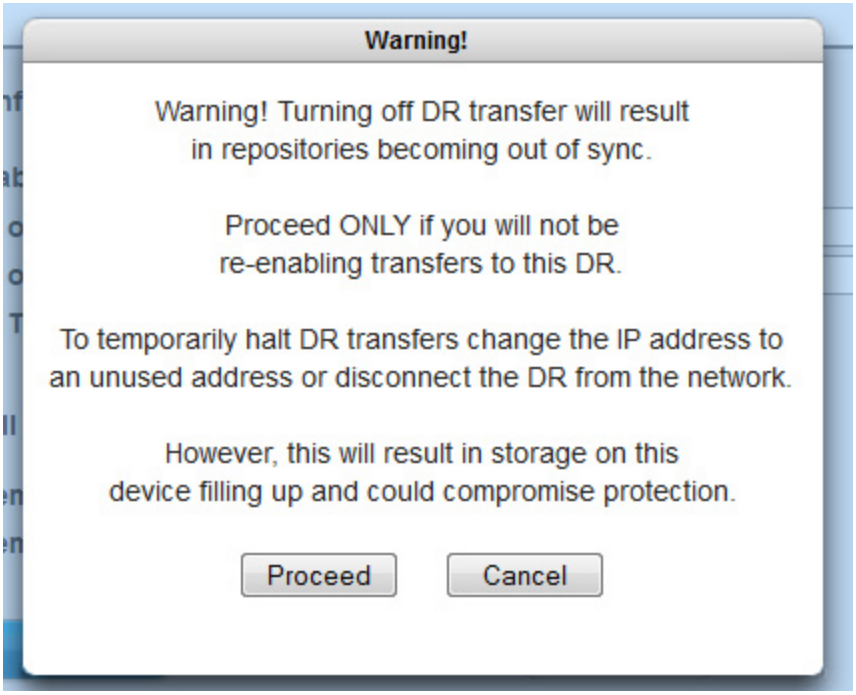
Set up trust relationships

To disable DR mirroring:

Before you disable DR mirroring, ensure that your onQ Appliance has enough disk space to accumulate backups. For the related alert, go to vDisk or Disk Space Alerts.

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.

6.13 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

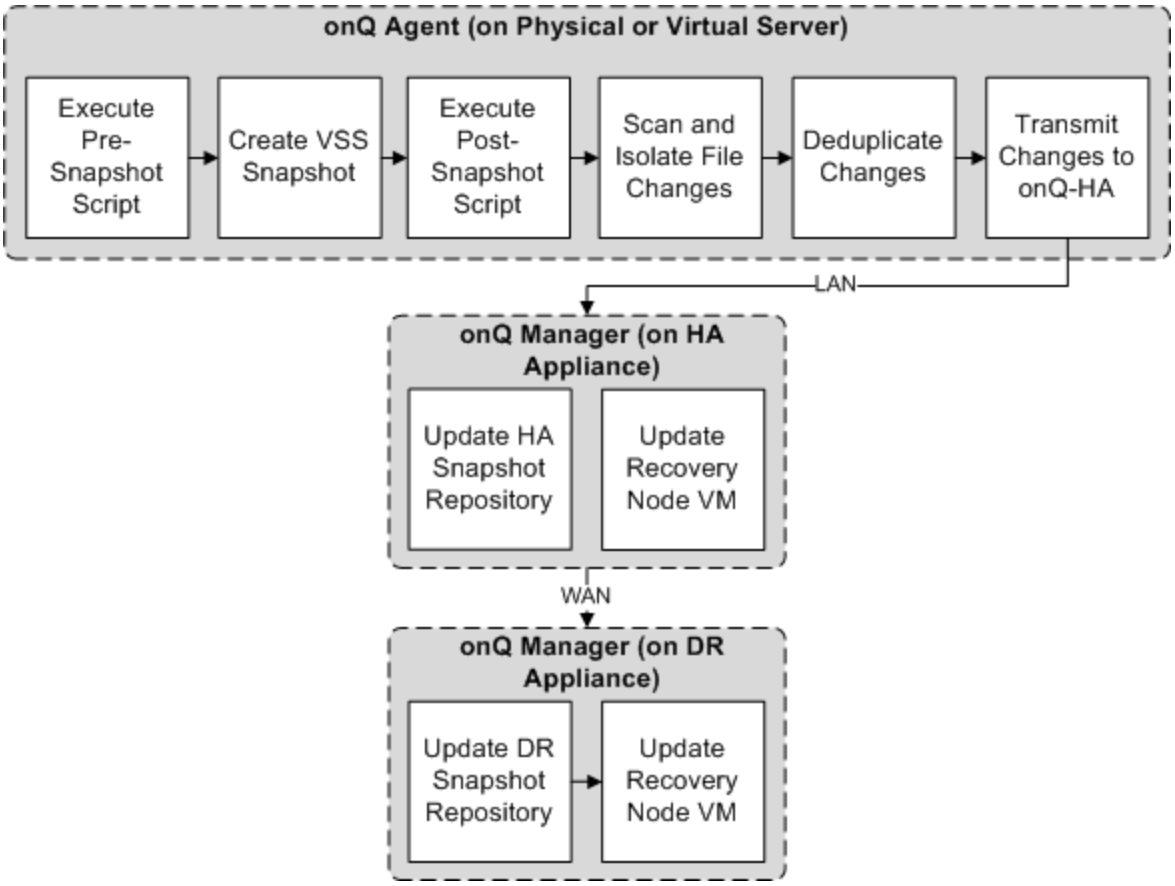
7

Backup and Restore

- Backup and Recovery Workflow
- Coordination of Backups and Recovery Nodes
- Edit RN services list
- Edit backup exclude list
- Edit backup include list
- Initiate immediate backups
- Export or import Repository
- 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
- Methods of Performing Incremental Backups
- Perform full rescan on PN
- Back up and restore Oracle 11g database

7.1 Backup and Recovery Workflow

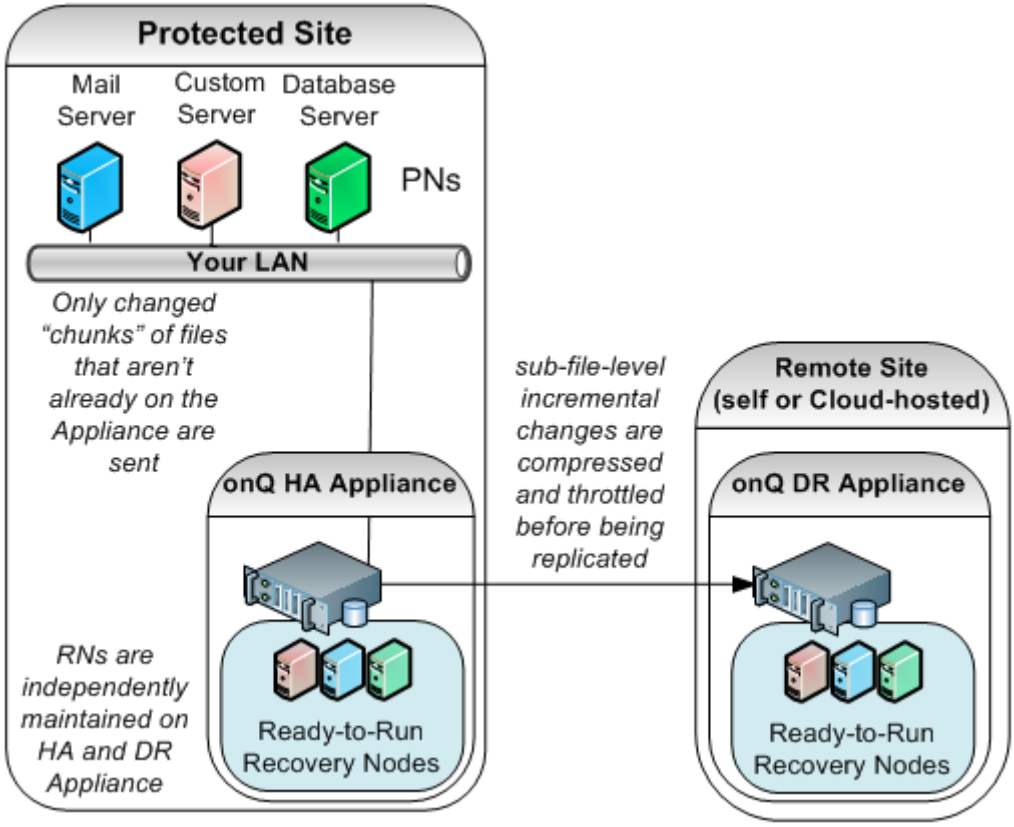
The repository is the deduplicated archive of all snapshots of all of your Protected Nodes. From the repository, onQ can reproduce a complete system image from any snapshot.



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 it detects a PN or backup failure, it issues an email alert. Your Administrator can then determine if the condition of the PN requires starting the HA or DR recovery nodes using the onQ Portal.



Related Topics

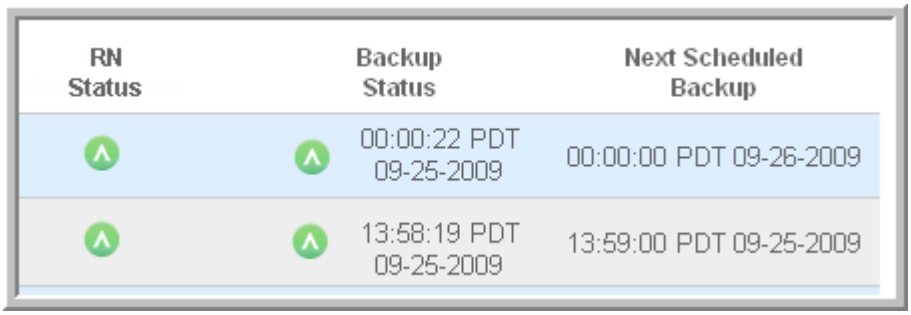
Coordination of Backups and Recovery Nodes
Initialize the DR Repository
Initiate immediate backups
Resize protected node's vdisk

7.2 Coordination of Backups and Recovery Nodes

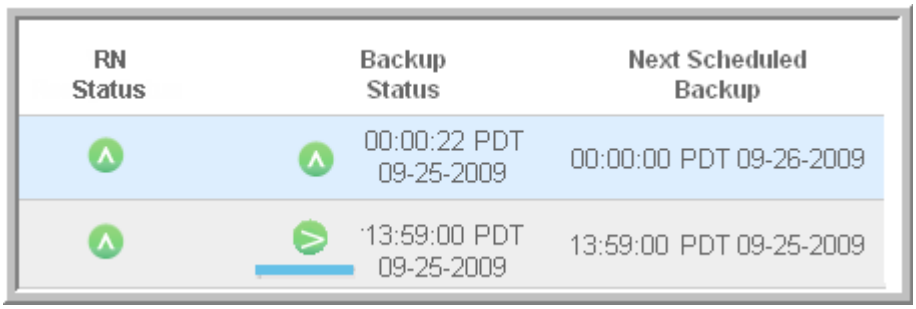
You can follow the HA Appliance's process of taking a backup and preparing a recovery node by watching the icons under **RN Status** and **Backup Status** on the **PROTECTED NODES** page.

The sequence is as follows:

1. Between backups, if protection is enabled, the icons under **RN Status** and **Backup Status** are each a green circle with an upward-pointing caret.

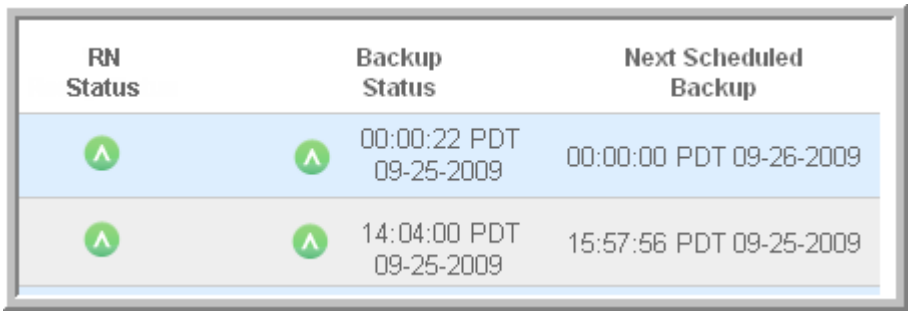


2. At the scheduled time, the backup begins. The caret on the green circle under **Backup Status** moves left to right, while a blue status bar underneath fills left to right.

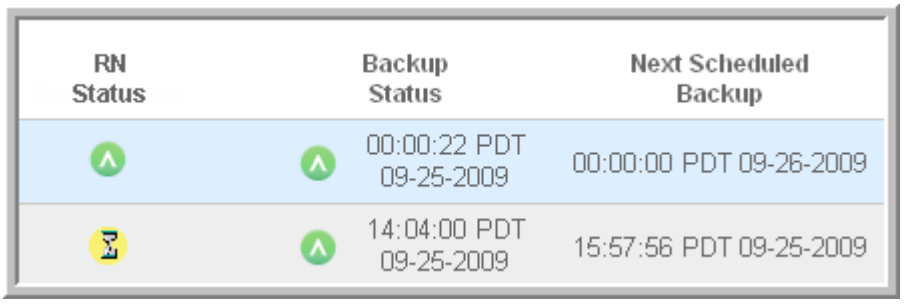


3. When the backup is finished, the dates under **Backup Status** and **Next Scheduled Backup** are updated and the circle under **Backup**

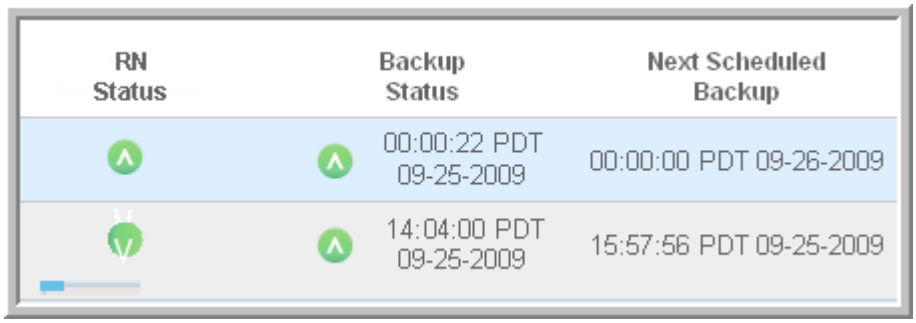
Status returns to a stable green circle with an upward-pointing caret.



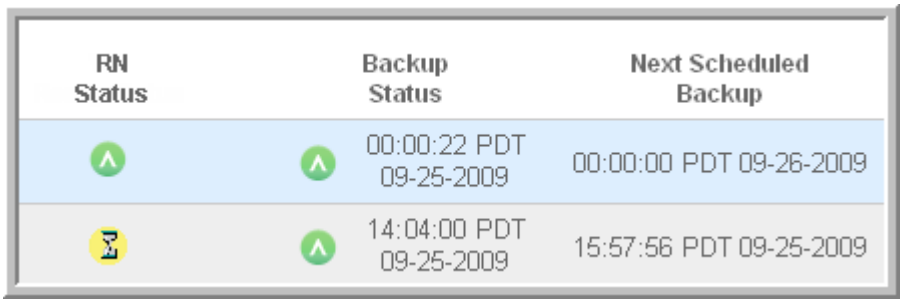
4. In a few moments, the circle under **RN Ready** changes to a yellow hourglass and remains so for about two minutes.



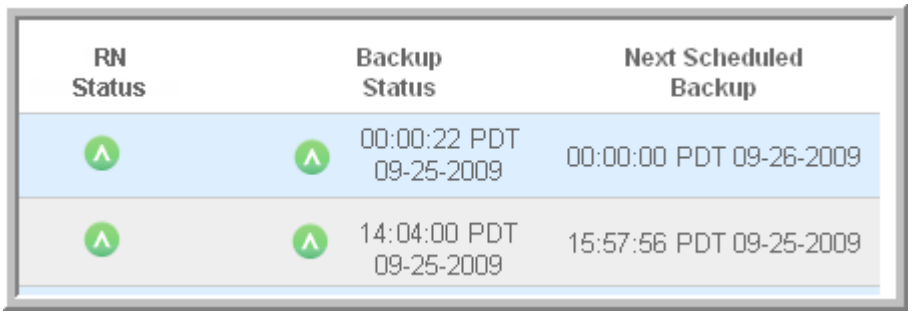
5. The icon then changes to a green circle with the caret appearing to fall top to bottom, filling a blue status bar beneath.



6. When the recovery node has been prepared, the circle again changes to a yellow hourglass and remains so for about one and a half minutes.



7. Finally the hourglass changes to a stable green circle with an upward-pointing caret. The circle remains until the cycle repeats at the next backup time.



Related Topics

Monitor protected nodes

7.3 Edit RN services list

An 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.

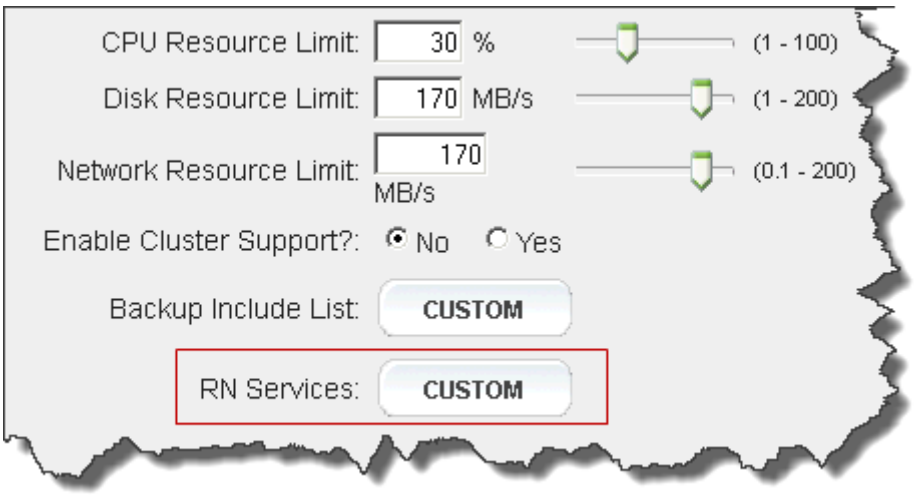
- If a custom list exists for an RN, the onQ Portal uses that list instead of the defaults list.
- If you later delete the PN, the onQ Portal deletes the custom list for that PN.

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** > **ADVANCED**.
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.

7.4 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.

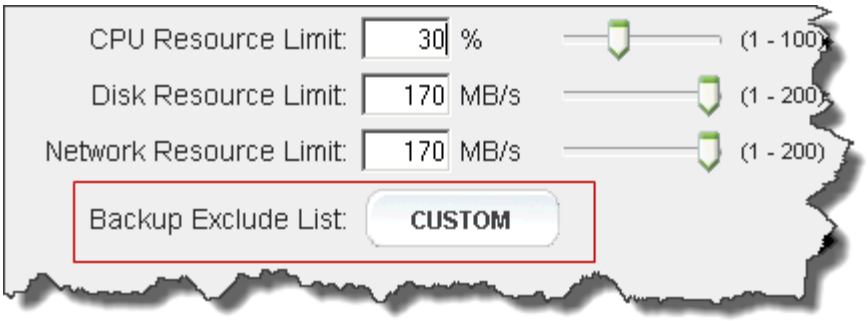
For an exclude list, Quorum recognizes file- and folder-names expressed as IEEE POSIX Basic Regular Expressions (BRE). You will find a synopsis of BRE syntax at http://en.wikipedia.org/wiki/Regular_expression.

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)

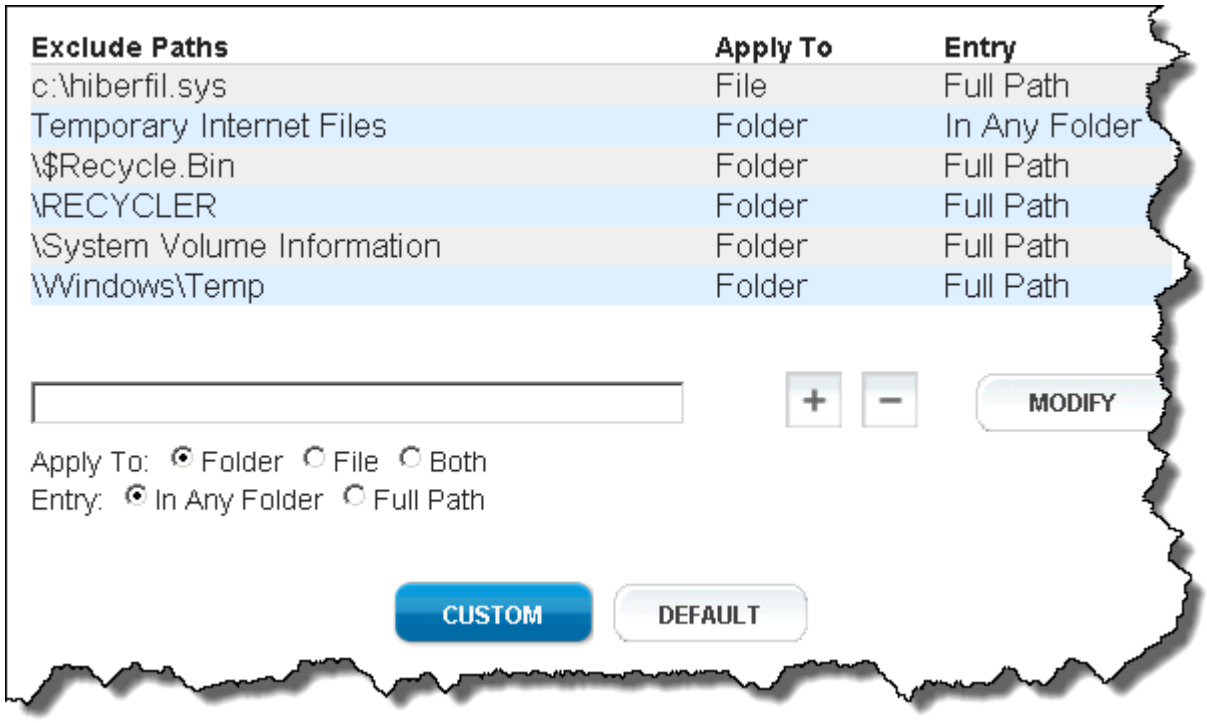
      & (ampersand)

   b. Select an **Apply To:** radio button that represents the path type. The path can be the folder or file or both.

   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 Exclude Paths, Apply To, and Entry choices mean "do not back up the contents" for `c:\hiberfil.sys`,

`Temporary Internet Files`, `\$Recycle Bin`, `\RECYCLER`, `\System Volume Information`, and `\Windows Temp`.



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.

7.5 Edit backup include list

Note: If you want file-only backups, you must manually enable this feature. Run the following on onQ: **touch /var/www/html/ENABLEBACKUPONLY**. Use the Include List feature outlined below to add the volumes, folders and files that you want onQ to back up.

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.

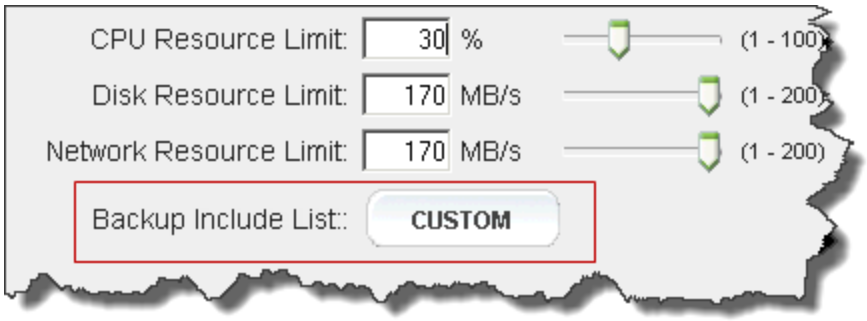
For an include list, Quorum recognizes file- and folder-names expressed as IEEE POSIX Basic Regular Expressions (BRE). You will find a synopsis of BRE syntax at http://en.wikipedia.org/wiki/Regular_expression.

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 Mode > ADVANCED**.

5. In the **Backup List** field, click **Custom.**



Note: If file only backup is set to **Yes**, the advanced menu now has the **Backup Include List** option.

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.

7.6 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
Initialize the DR Repository
Resize protected node's vdisk

7.7 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. 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 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 Appliances, it might be more efficient to export onQ's deduplicated snapshot storage (or "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 protectionon the HA Appliance.

Warning: Do not restart 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. Hover-over the RN status for current state information.

   Alternatively, you can stop the backups as outlined in "Stop in-progress backups" on page 108.
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 protectionon the DR Appliance.

4. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **REPOSITORY** page.
5. Click **Import**.
6. Wait for import to complete.
7. Restart protection on HA Appliance.
8. Restart protection on DR Appliance.

Related Topics

Initial Backup
Initialize the DR Repository

7.8 Delete snapshots

When you delete a node, you might also want to delete its snapshots. It is not possible for *you* to delete snapshots at this time. Contact Quorum Support at support@quorum.net so that we can delete the snapshots for you.

7.9 Stop in-progress backups

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.

7.10 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.

7.10.1 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 Ontrack PowerControls.

To retrieve a snapshot, using WSR:

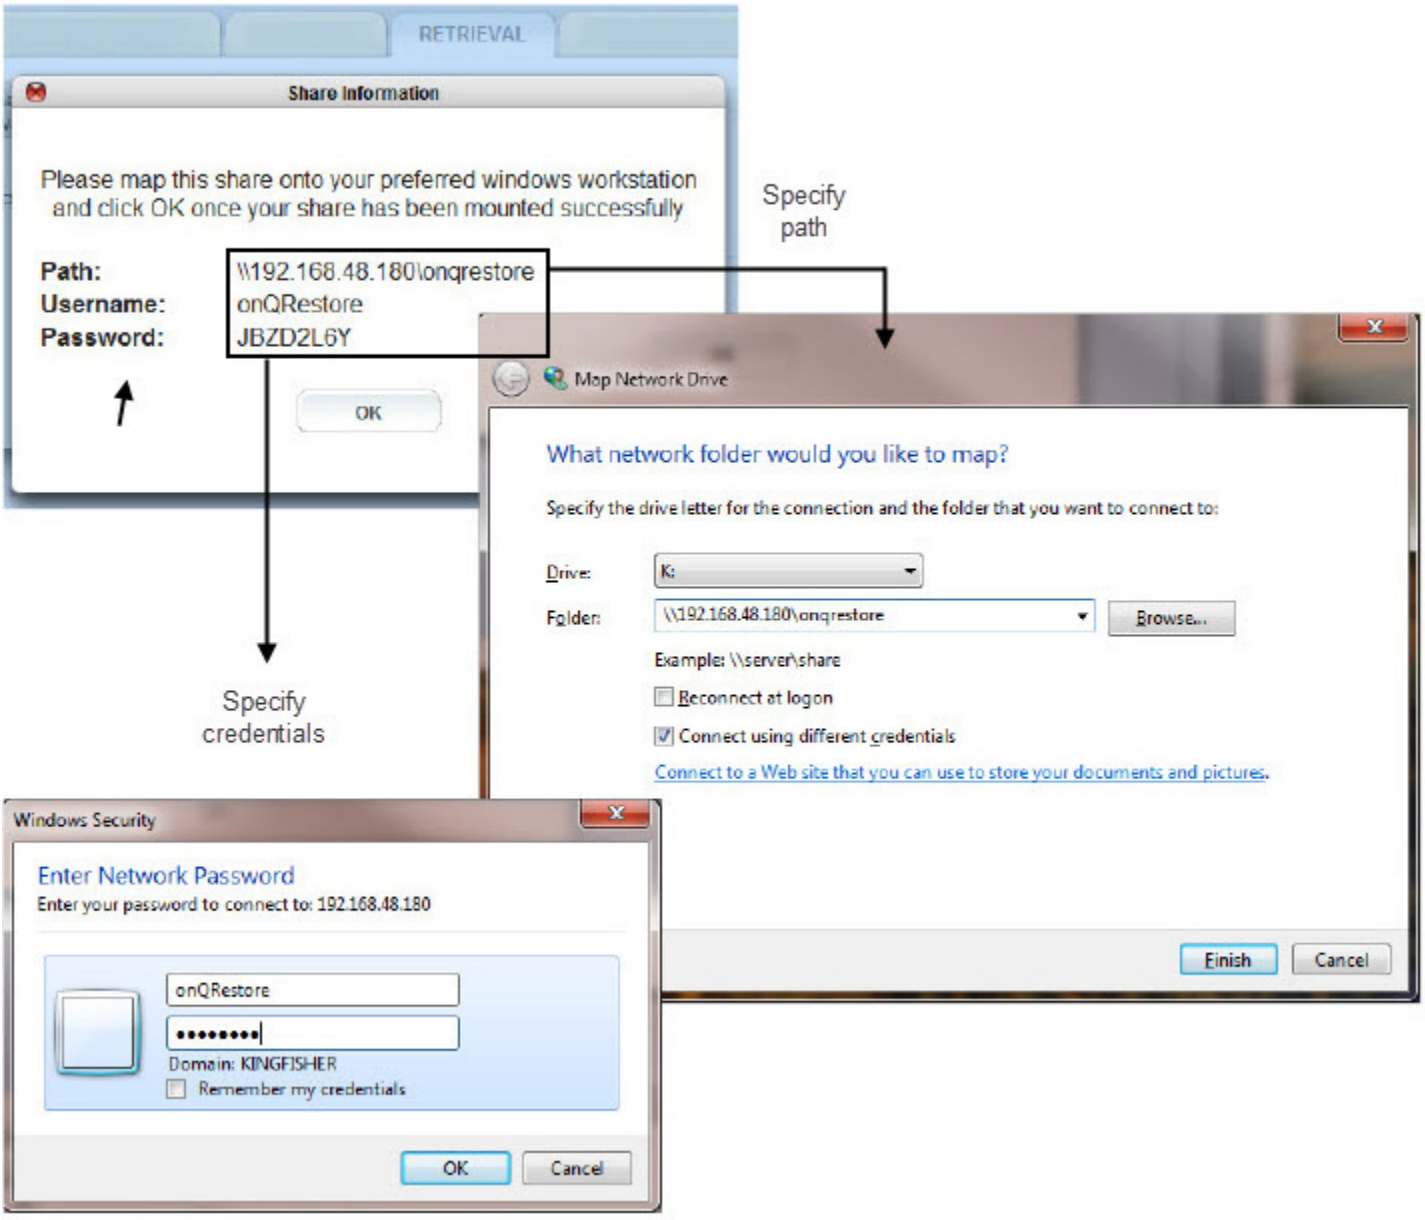
1. Log on to the Appliance's onQ Portal.
2. Click on the **RESTORE** tab > **WINDOWS SHARE** page. onQ displays an empty page if no share was activated.
3. Create a share, then map to this share:
   a. Click on the **ACTIVATE SHARE** button.

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**.

7.10.2 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).

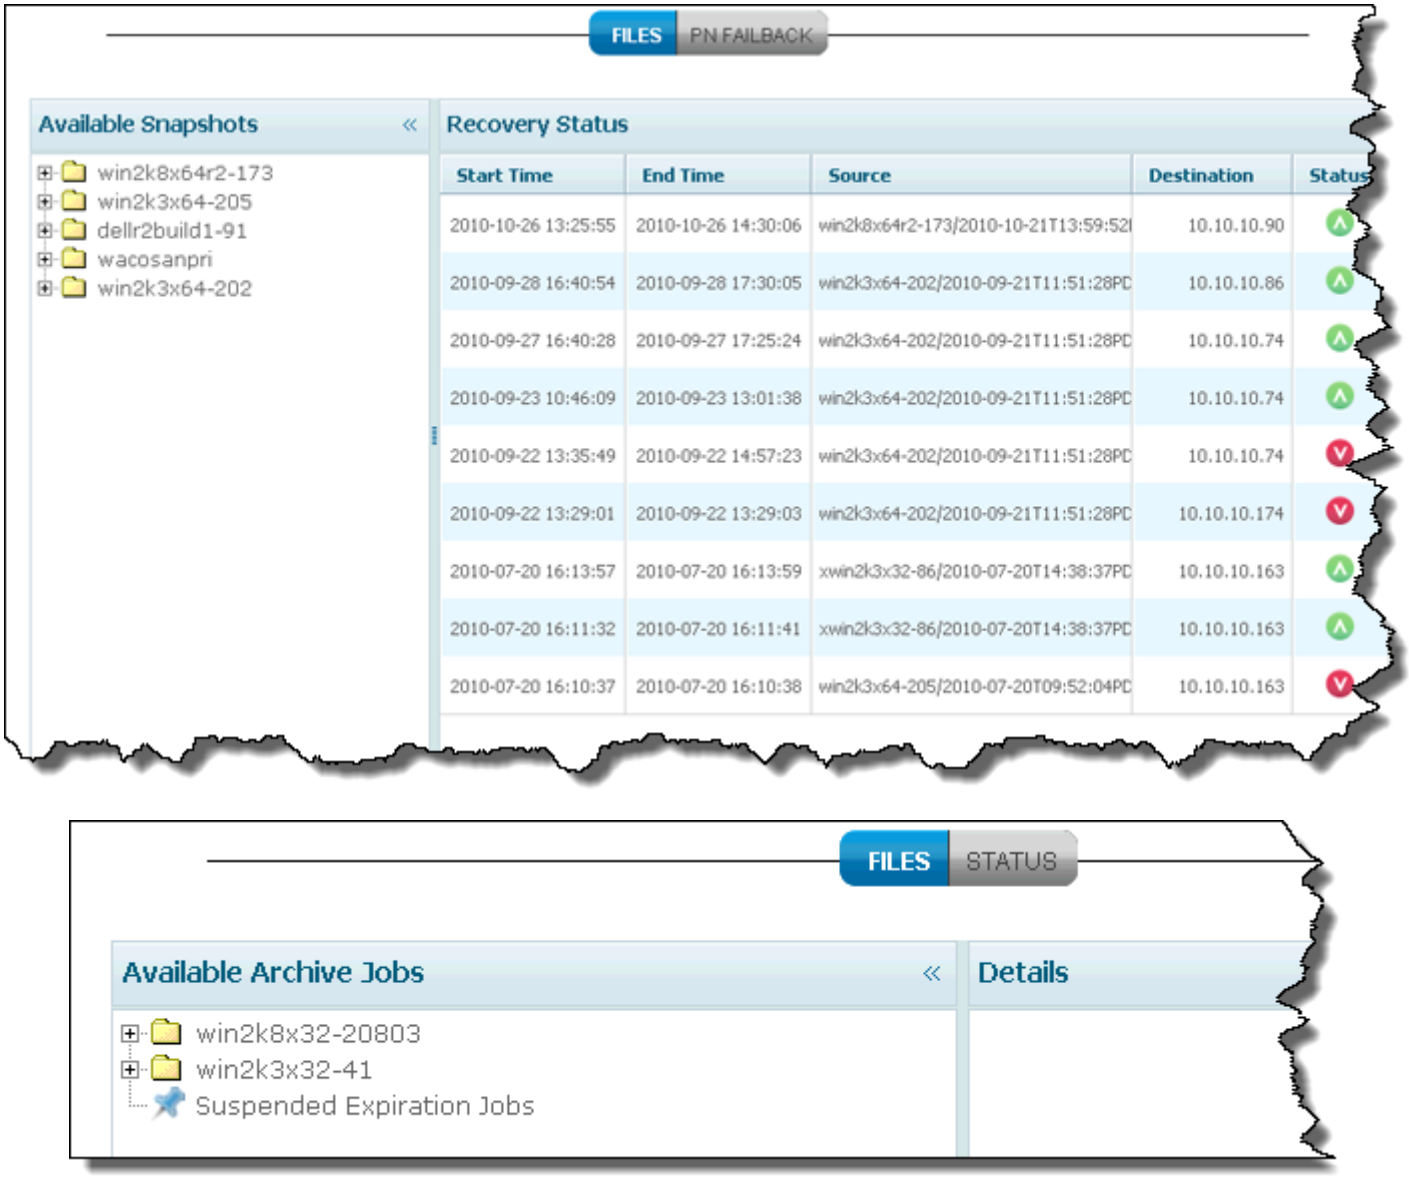
If, however, your protected node's server failed, you need to perform a bare metal restore (BMR) as outlined in "Perform Bare Metal Restore" on page 134. You can choose to *replace* the current file, folder, or drive, or to *store* it on any protected node or standalone server so long as it is reachable 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:

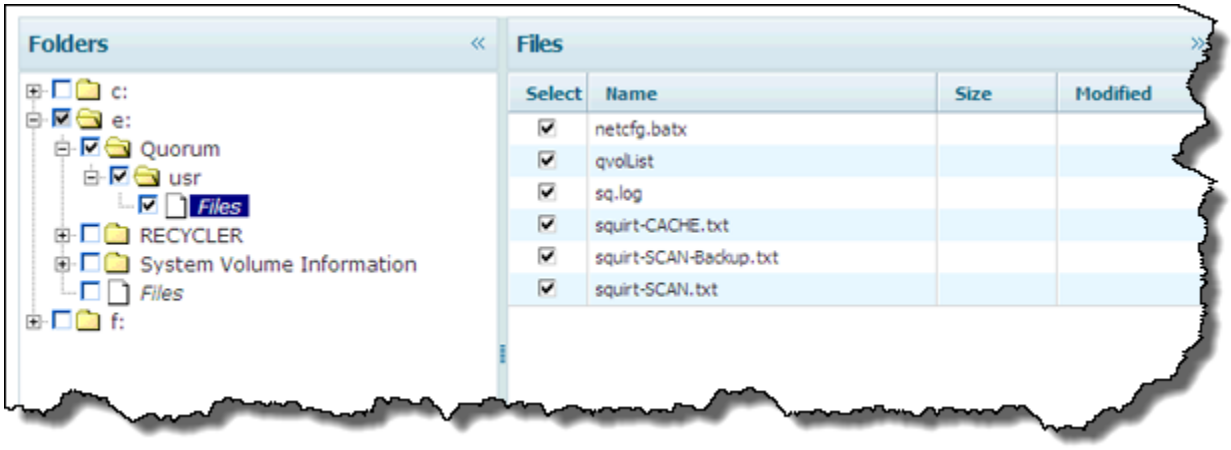
1. Log on to the Appliance's onQ Portal.
2. Click on the **RESTORE** tab > **FILES** page.

A page similar to the following appears:



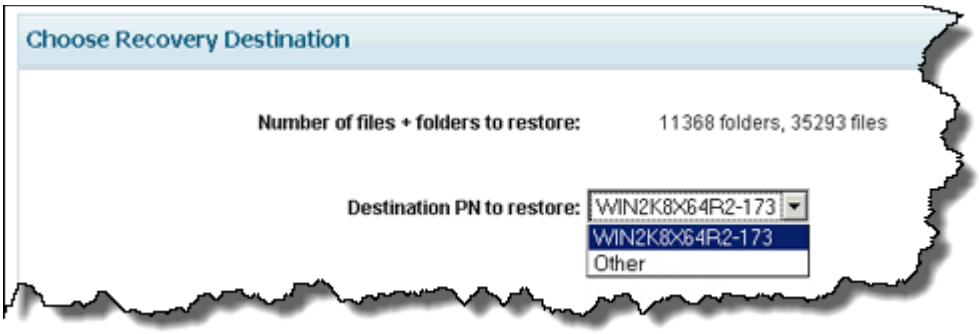
3. In the **Available Snapshots** pane, expand the folder for the protected node.
4. Double-click on the snapshot that contains the data you want to restore.

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.
6. Before you proceed with the restore, identify the host to which you want to restore the files.
   - **You want to restore to a Protected Node**. Do so if you want to restore to the same or other Protected Node.
   - **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. Based on your decision in Step 6, in the **Destination PN to restore** drop-down, specify the recovery destination — the host where you want to store the retrieved files. The files will automatically be saved to:
   - (Windows) - `C:\onQRestore`.
   - (Linux) - *mountpoint*`/onQRestore`. For example, `/onQRestore`

if you chose / or /boot/onQRestore if you chose /boot.



Note: 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.

The destination can be on the same PN, any other PN, or some other destination.

- (**PN hostname**) Choose this option if you chose "Protected Node" in Step 6, then select the drive. **Skip** the next step.
- (**Other**) Choose this option if you chose "Standalone Server" in Step 6. When you select **Other**, a text box appears labeled **Destination hostname or IP address**. Type the address of your desired destination server. **Proceed** to the next step.

8. If you want to restore to a standalone server, which, unlike a protected node does not already have the onQ Monitor installed, do the following to open up a communication "channel" between the onQ Appliance and the standalone server:

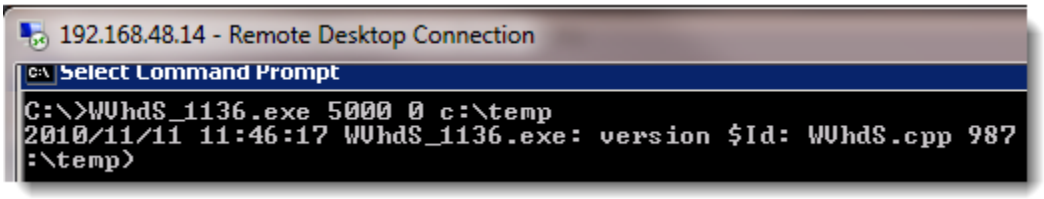
   a. Download and execute one of the following files on that remote host, depending on your operating system.

   - **(Windows)** WVhdS. You need only download and execute this file on the remote host as the example below shows. 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)

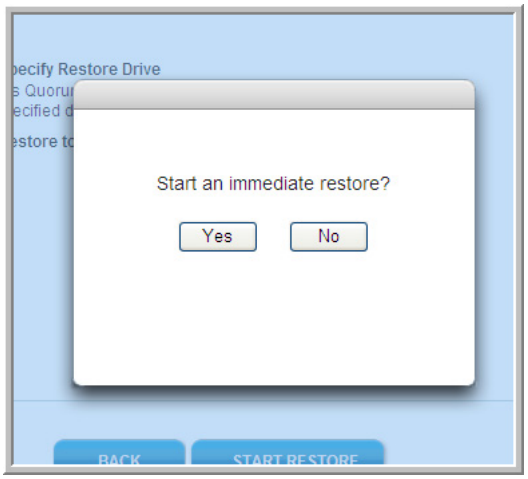


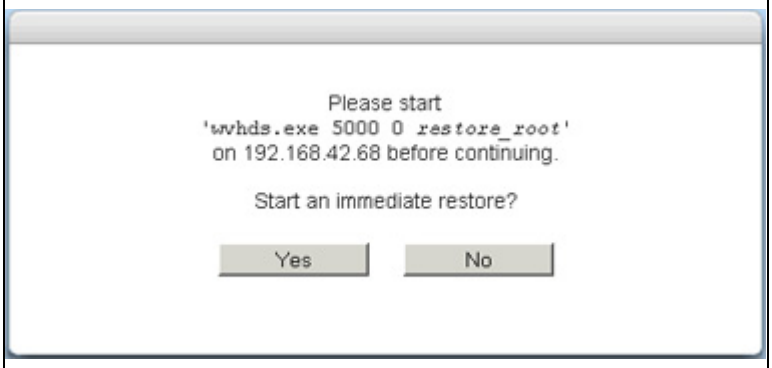
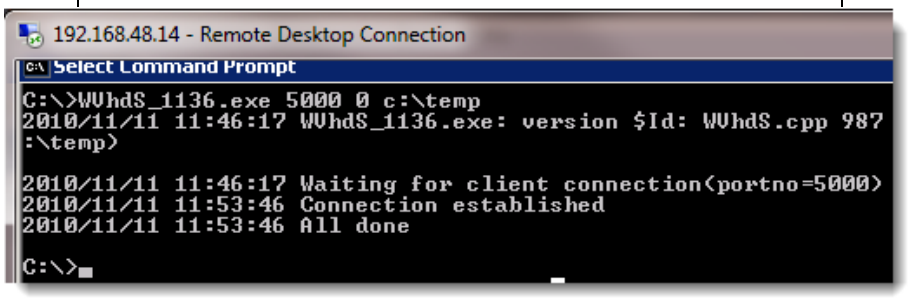
- **(Linux)** `onQFLRSelf.sh`. Automatically untars a file to the recipient that you define:

```
# cd <restore-location>
# wget http://<OnQ-IP-address>/onQFLRSelf.sh
# sh onQFLRSelf.sh
```

9. 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 **PN host-name** | **Specify Restore Drive**:	

<table>
<tr>
<td rowspan="2">If you chose PN host-name | Specify Re-store Drive
OR
Other:</td>
<td>
</td>
</tr>
<tr>
<td>If restore completes, output indicates All done:

</td>
</tr>
</table>

Afterward, the onQ Portal returns you to the main **RESTORE** page.

10. Verify that your restore completed successfully:

    a. Observe Snapshot status:
    - 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 past 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 failure.
	Incomplete(ReplyError)  carat is red and points downward	Restore failed. If you roll your cursor over the icon, you see a reason for failure.

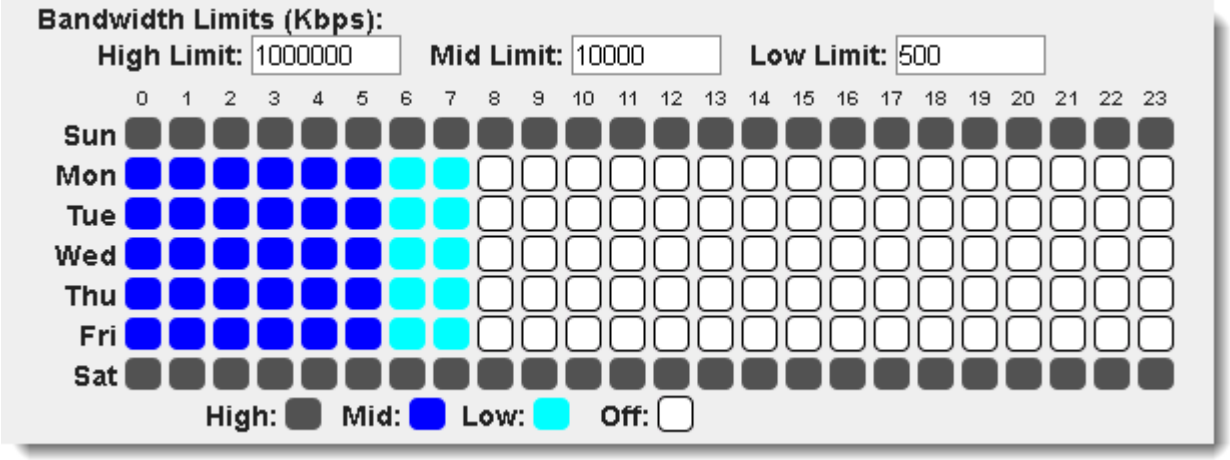
b. (Standalone Server) Verify that the restored files appear in the expected destination.

7.11 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 processes place too much of a burden on your network.

The onQ Portal enables you to create daily schedules that adjust the bandwidth to allow high utilization during non-working hours and low utilization during working hours:



To adjust bandwidth throttling between HA Appliance and DR Appliance:

1. Log on to either 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 to the select the limit for each day and hour, then **SAVE**.

To adjust bandwidth throttling between DR Appliance and DR Mirror:

1. Log on to either 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 to the select the limit for each day and hour, then **SAVE**.

Related Topics

Limit resource utilization on PNs

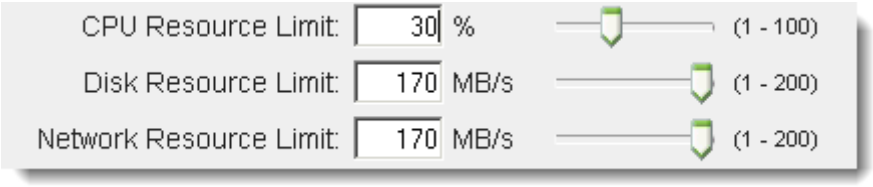
7.12 Limit resource utilization on PNs

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

7.13 Run custom backup scripts

onQ supports pre-snap and post-snap scripts to enable you to quiesce an application or database for backup without manual intervention.

- 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.

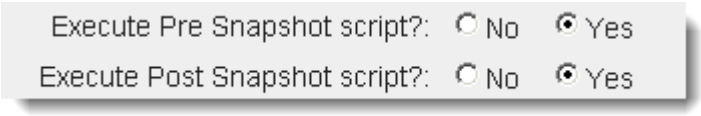
You can do so on a per-protected-node basis by saving your scripts to the PN, then registering them with the HA Appliance.

If you want to run these same scripts on all protected nodes, you must repeat the procedure below for each protected node.

To enable custom backup scripts:

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 a `.bat` example scripts that 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.
   - **`pre_snap.sh`** and **`post_snap.sh`**: For Linux, these scripts were saved to the destination you chose during the installation. The default location is `/opt/quorum/bin`.
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`.



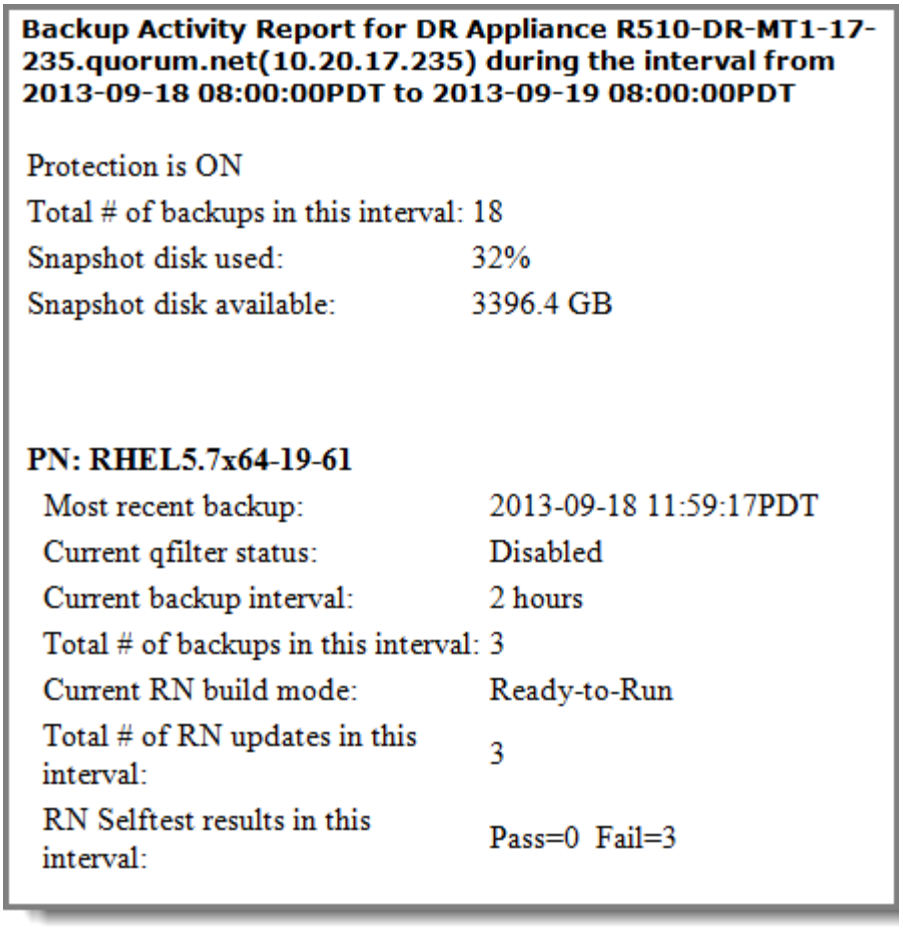
Related Topics

Back up and restore Oracle 11g database

7.14 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. The report also provides statistics on the amount of disk space that remains on the Vdisk and status of each PN's Self Test.



Backup Activity Report for DR Appliance R510-DR-MT1-17-235.quorum.net(10.20.17.235) during the interval from 2013-09-18 08:00:00PDT to 2013-09-19 08:00:00PDT

Protection is ON

Total # of backups in this interval:	18
Snapshot disk used:	32%
Snapshot disk available:	3396.4 GB

PN: RHEL5.7x64-19-61

Most recent backup:	2013-09-18 11:59:17PDT
Current qfilter status:	Disabled
Current backup interval:	2 hours
Total # of backups in this interval:	3
Current RN build mode:	Ready-to-Run
Total # of RN updates in this interval:	3
RN Selftest results in this interval:	Pass=0 Fail=3

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.

Related Topics

Resize protected node's vdisk
Interpret e-mail alerts
Modify e-mail alert settings

7.15 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.

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.

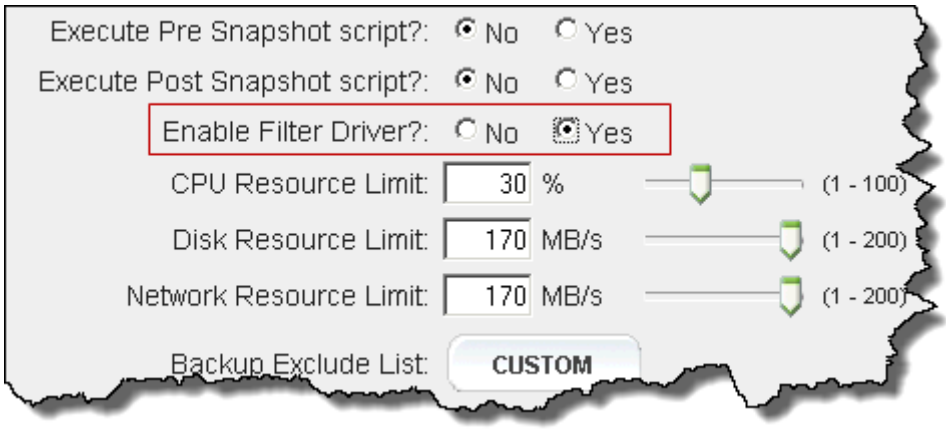
Though you don't need to know the "ins and outs" of each solution, you do need to understand when each performs optimally, so that you can choose

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

7.16 Perform full rescan on PN

In most cases, you will not need to perform a full rescan. However, you might want to 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**.

7.17 Back up and restore Oracle 11g database

The following procedure walks you through how to restore an Oracle 11g database on a Linux Recovery Node using example scripts and the Oracle Tool Kit. These scripts execute a hot backup of your Oracle database using `RMAN` and then recover the Oracle database onto the RN.

To install the RMAN scripts:

This procedure works for Linux 6.2 RNs, though this procedure should also work on other os versions.

1. Install Oracle Tool Kit on the PN.

   This procedure works with version 1_0_2_1_5, though this procedure should also work on later OTK versions. For this version, OTK was installed in the following locations:
   - `/opt/oracle`
   - `/opt/oracle/admin/dev`
   - `/opt/oracle/otk`
2. Download the example scripts.
3. On the Linux PN that's running Oracle 11g database, copy the following scripts to `/opt/quorum/bin`:
   - `pre_snap.sh`
   - `post_snap.sh`
   - `armutil.sh`
   - `ora_00_config.sh`
   - `ora_02_catalog.sh`
   - `ora_05_restore.sh`
   - `ora_01_util.sh`
   - `ora_03_pre.sh`

Note: The comments at top of each script explain the purpose of each script.

4. Copy the `armSender` and `vmmonCtl` from x64 or i386 folder (depending upon the PN's architecture) to `/opt/quorum/bin` on the Linux PN that's running Oracle 11g database.

The `armSender` and `vmmonCtl` are tools to communicate status to the onQ Manager.

5. Log on to the HA Appliance's onQ Portal.
6. Stop protection.
7. From the **PROTECTION CONFIG** tab, select the PN that's running Oracle 11g database, then **MODIFY** > **ADVANCED**.
8. Enable the pre snapshot scripts on that PN by selecting the **Execute Pre Snapshot script?: Yes** radio button.

To test the scripts:

1. Load some known data into the Oracle database (for example in DEV catalog).
2. Perform backup of the PN.

   The Event log displays the following messages as onQ builds the RN:
   - Pre-snapshot command returned: rman backup succeeded (as coded in pre_snap.sh), if rman backup passes.
   - Pre-snapshot command returned: rman backup failed, see /tmp/ora_03_pre.out (as coded in pre_snap.sh), if rman backup fails.
3. Restore the database using the procedure below.

To restore an Oracle 11g 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 production mode, after the RN build completes.
2. On the RN, run the recovery script as Oracle user to restore the Oracle database onto the RN:

```
# su - oracle
# ora_05_restore.sh
```

3. Verify that the RN's Oracle database shows the expected data.

Related Topics

Run custom backup scripts

8

Bare Metal Restore

- Perform Bare Metal Restore
- BMR Errors

8.1 Perform Bare Metal Restore

Before you begin, learn more about BMR support and limitations:

- "Limitations" in onQ Release Notes
- "Support" in onQ Release Notes

A Bare Metal Restore (BMR) refers to the process of recovering from a server failure by restoring a Recovery PN from a snapshot on the onQ Appliance.

To perform this BMR, use the Quorum Ultimate Automated Recovery Kit (QUARK), which has an easy-to-use wizard that walks you through the entire BMR.

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.

To perform a bare metal restore:

In this procedure, *Recovery PN* (also called *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, this procedure assumes that you have a snapshot of your virtual machine or a template from which to deploy a clone. In any case, this virtual machine is only a "shell" onto which QUARK can load its snapshot(s). 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.

1. Verify that your platform is supported. Go to "Support" in onQ Release Notes.
2. If you have an EFI partition, make it bootable.
3. Make BIOS changes and launch QUARK wizard:
   a. 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.

   b. Configure the Recovery PN's BIOS to boot QUARK.
   c. 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, the **Add Drivers** page appears.

Note: To manually launch this **Add Drivers** utility, run the following command from the command prompt:

```
# QUARK_Wizard.exe AddDriver.
```

5. (If necessary) In the **Add Drivers** page, upload the missing network driver, then **Next**.
6. From the **Continue onQ Restore** page, select the **Continue with onQ Restore** workflow, then **Next**.
   - If the wizard detects a DHCP server, the wizard displays the assigned default IP address and netmask in the **Network Configuration** page.
   - If DHCP is not enabled on the network, the wizard displays the assigned default Windows workgroup address.
7. (Optional) In the **Network Configuration** page, change the default network configuration, then **Next**. The **Log on to onQ Appliance** page appears.

   Select the **Set Static IP** check box to set the default IP address to a static IP.
8. In the **Log on to onQ Appliance** 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 analyses 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 **PN Selection** page.

   If this process fails, browse the list of error messages in BMR Errors to correct the problem.
9. 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**.

   If the onQ Appliance doesn't have the incremental backups that are needed to perform the restore, the wizard displays the **Full BMR**

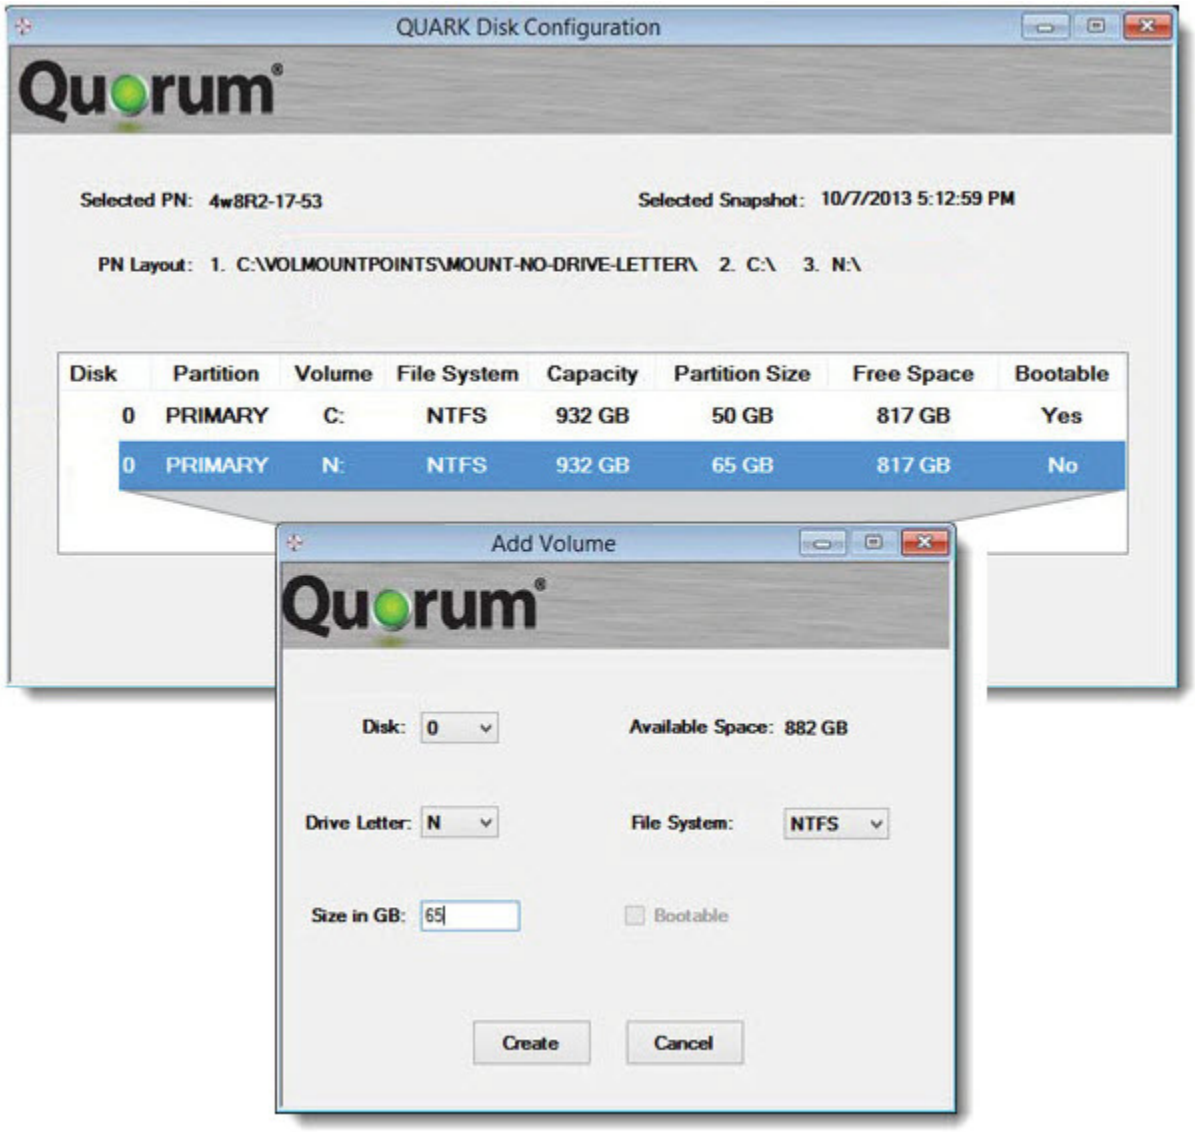
option only; otherwise, both options are available. If you perform a Full BMR, the **Select Snapshot** page appears.

10. In the **Select Snapshot** page, select the snapshot to which you want to restore the Recovery PN, then **Next**.

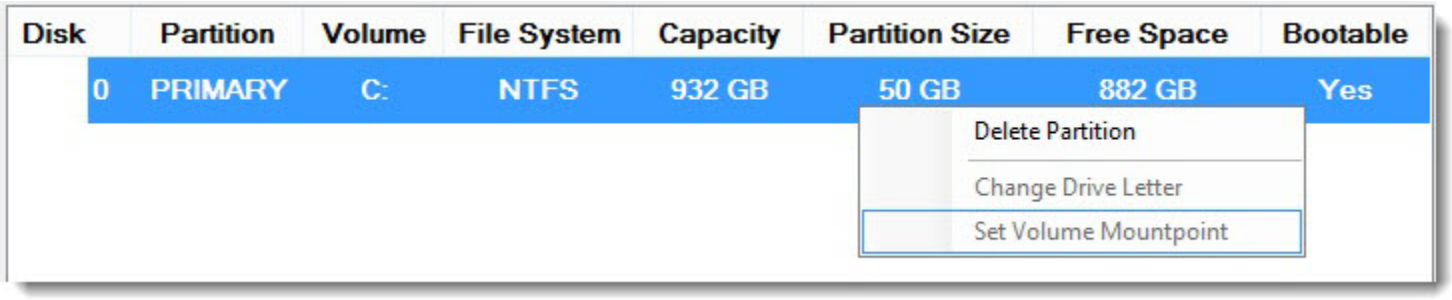
    The **QUARK Disk Configuration** wizard appears.

11. Configure the Recovery PN with the same disk configuration as the original PN:

    a. For every volume that appears in the PN Layout field, add a corresponding volume. You can change the size of the drive.

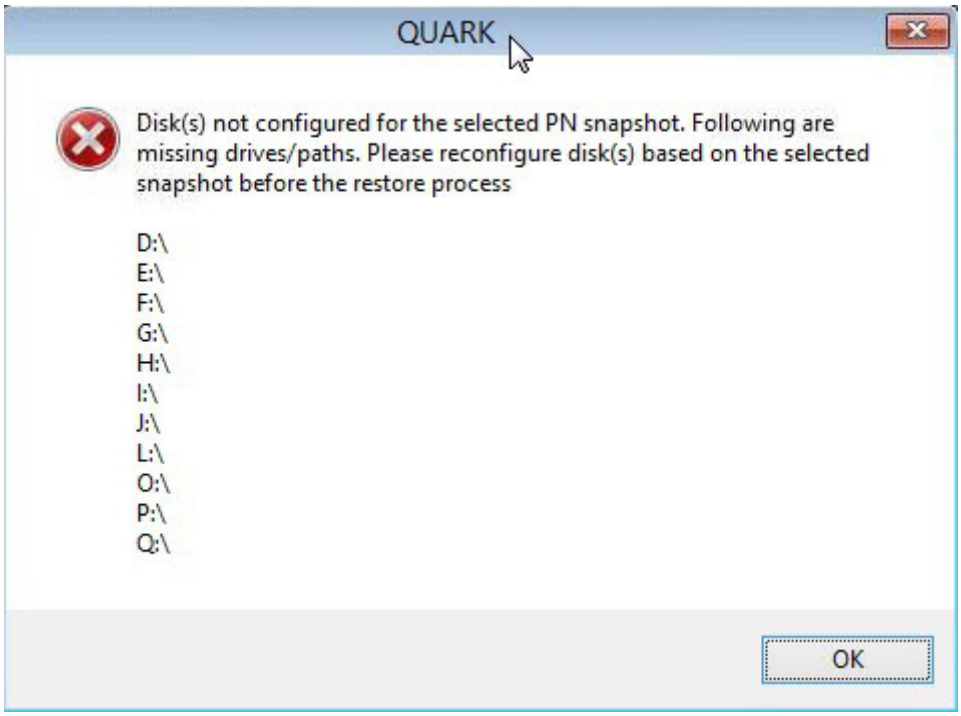


Moreover, if you want to delete a volume, right-click on the volume and **Delete Partition**.



b. When the **Do you want to process with the restore?** dialog appears, click **Yes**.

If you receive the following error message, the volumes that you added do not match the PN layout. You must change the configuration to match the PN layout.



c. 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 Failback Progress** page appears.

12. Wait from the onQ Appliance to restore the Recovery PN to the snapshot you selected, click **Next**.

    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 BMR Errors to correct the problem.

13. Do the following:

    a. (If necessary) From the **P2P Adjust OS Wizard**, select the Recovery PN's operating system, 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.

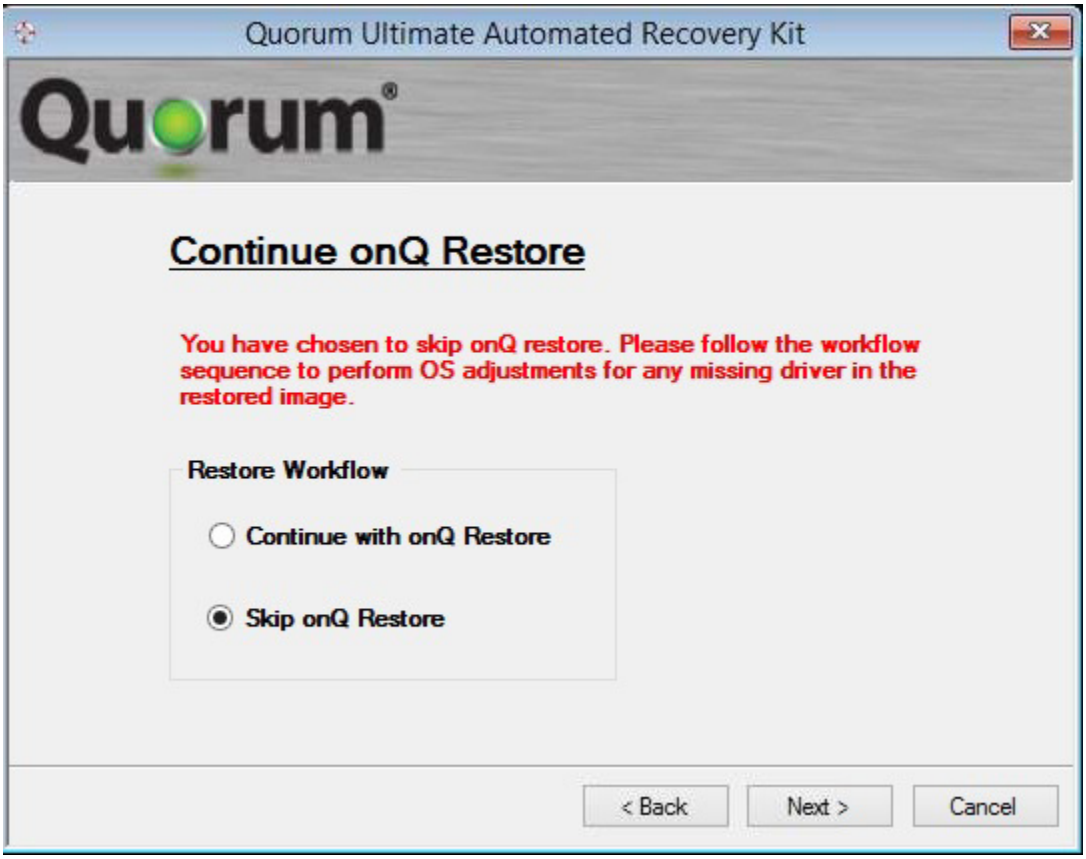
    b. Run one of the appropriate batch files from the `B:\BCV\Quark` folder.

Windows 2012 Server, 64-bit	`INJ_Win2012_BT.BAT`
Windows Server 2008 R2, 64-bit	`INJ_2K8_R2BT.BAT`
Windows Server 2008, 32-bit	`INJ_2K8_32BT.BAT`
Windows Server 2008, 64-bit	`INJ_2K8_64BT.BAT`
Windows 7, 64-bit	`INJ_2K8_R2BT.BAT`

You must do so now while the QUARK environment is still active and before the restore target reboots. Your restore target might not boot if you perform this step out of sequence.

c. Select **Reboot the system** or **Shutdown the system** to complete the BMR, then **Finish**.

If you reboot without adjusting for new hardware, you can always perform the adjustments after the restore completes. Simply launch the wizard again, by repeating Step 4, then select the **Skip onQ Restore** workflow from the **Continue onQ Restore** page.



8.2 BMR Errors

Error	Cause	What do I do?
`Authentication failed for onQ`	You typed the incorrect credentials.	Try again.
`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.
`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.

Related Topics

"Bare Metal Restore Support" in onQ Release Notes

Perform Bare Metal Restore

Perform File-level Restore

9

Security and Communications

- Test HA-to-DR Link
- Set up trust relationships
- Modify TCP Ports for HA-DR Communications
- Synchronize system time
- Secure HA connection to PNs
- Restart the onQ Service

9.1 Test HA-to-DR Link

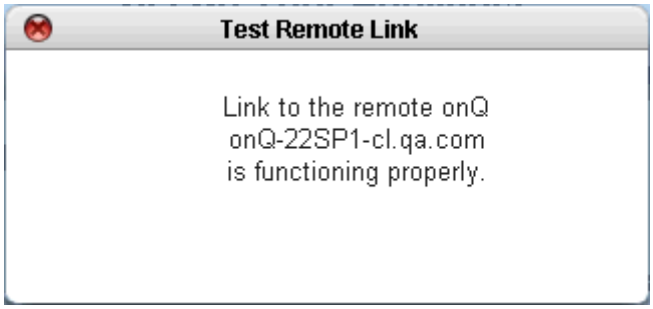
Use this procedure to:

- 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

9.2 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 Fail over HA Appliance to DR Appliance).

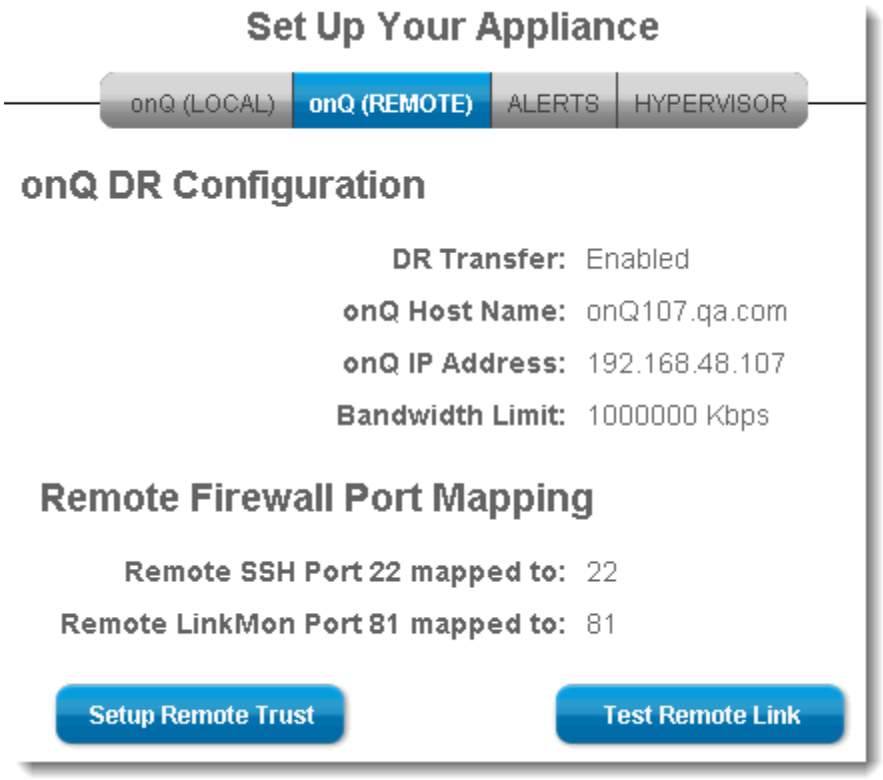
Therefore, just in case you have a disaster, it's best practice to set up the trust relationship in advance so that it's immediately available when you need it. As such, consider performing the following procedure on both the HA Appliance and, if you have a DR Mirror, the DR Appliance.

To set up the trust relationship between HA Appliance and DR Appliance:

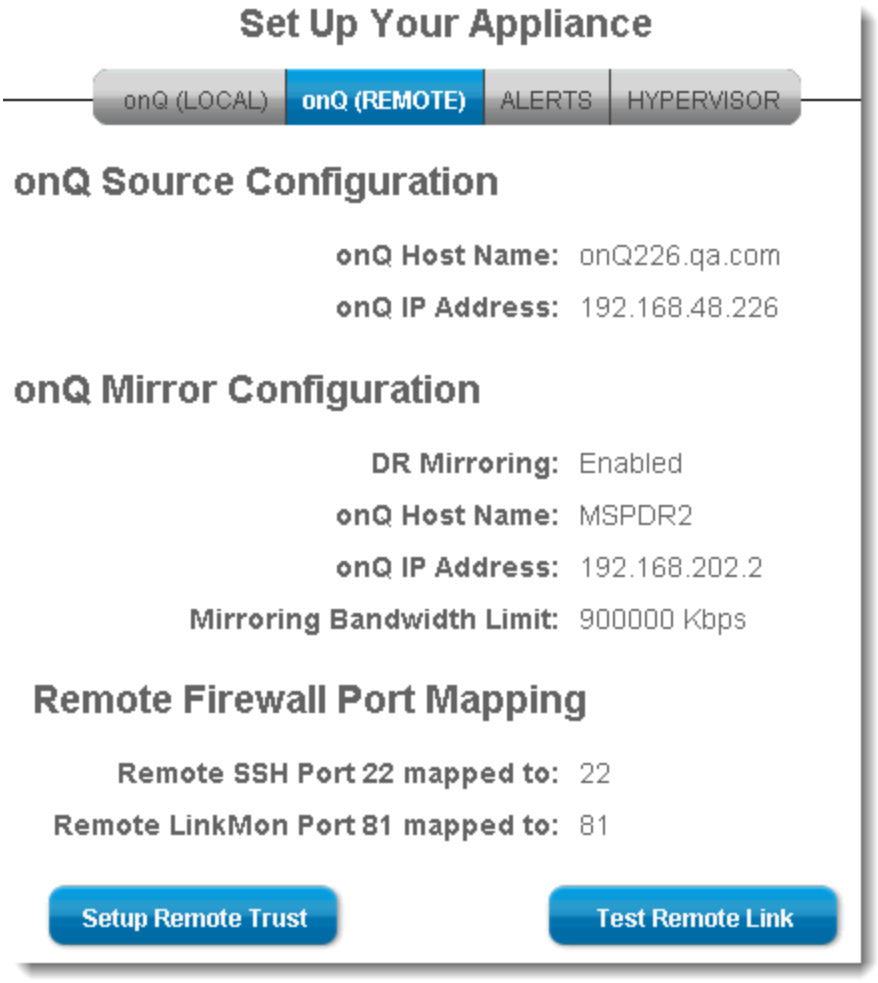
1. Log on to the HA Appliance's onQ Portal.
2. Go to **APPLIANCE CONFIG** tab > **onQ (REMOTE)** page.

   Simply click on the **Set Up Remote Trust** button.

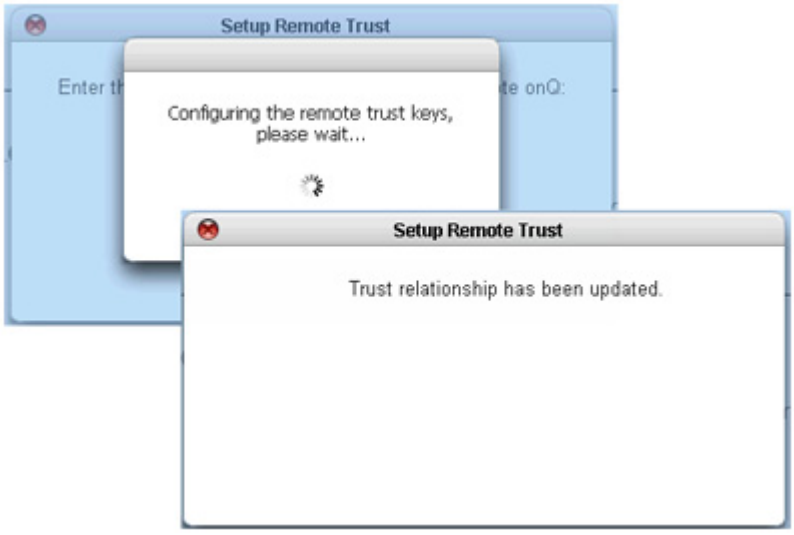
Example of HA Appliance:



Example of DR Appliance:



3. Type the VARAdmin password for the *other* onQ Appliance, then **Update Trust**.

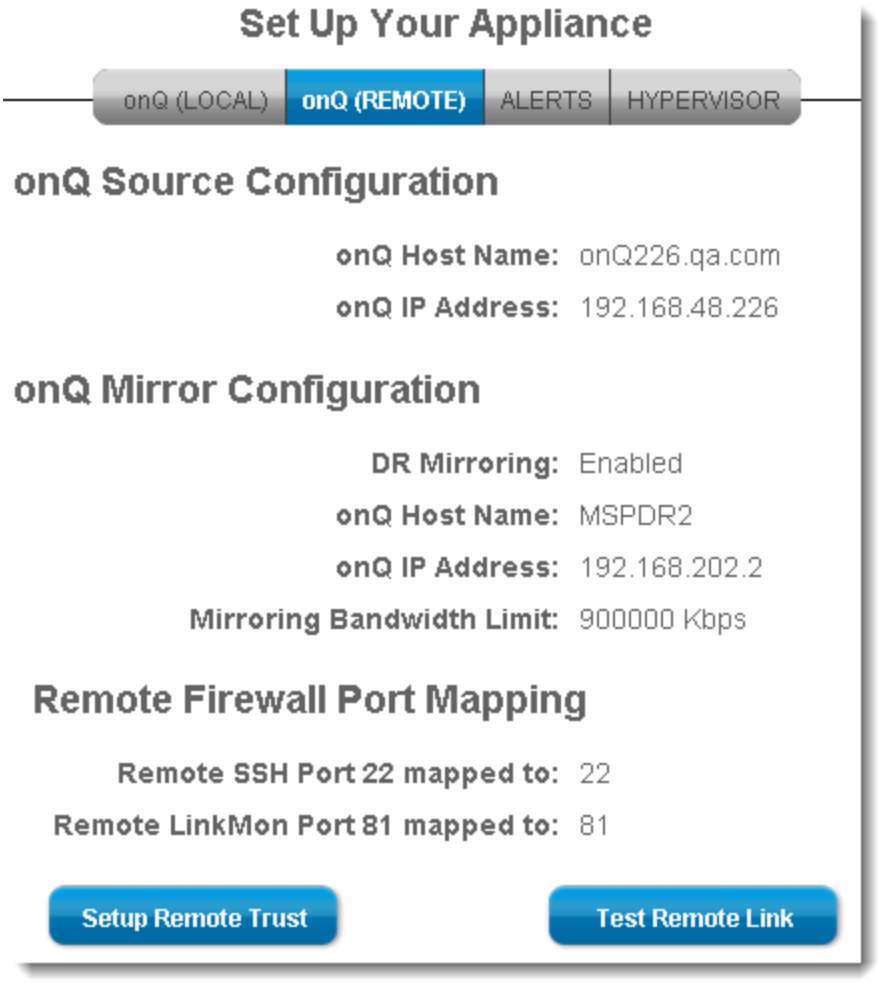


4. Test the link.

To set up the trust relationship between DR Appliance and DR Mirror:

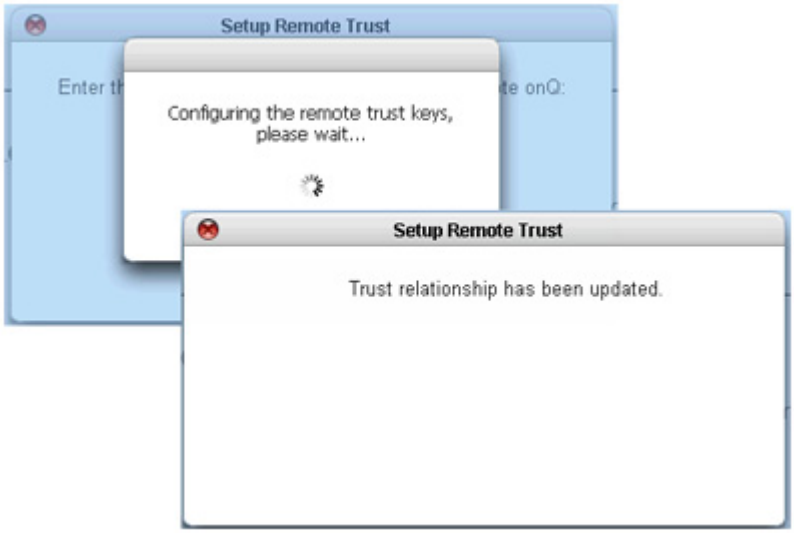
1. Log on to the DR Appliance's onQ Portal.
2. Go to **APPLIANCE CONFIG** tab > **onQ (REMOTE)** page.

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.

9.3 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.

9.4 Synchronize system time

The HA Appliance uses an NTP Server for its system time.

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 Troubleshoot

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.

9.5 Secure HA 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.

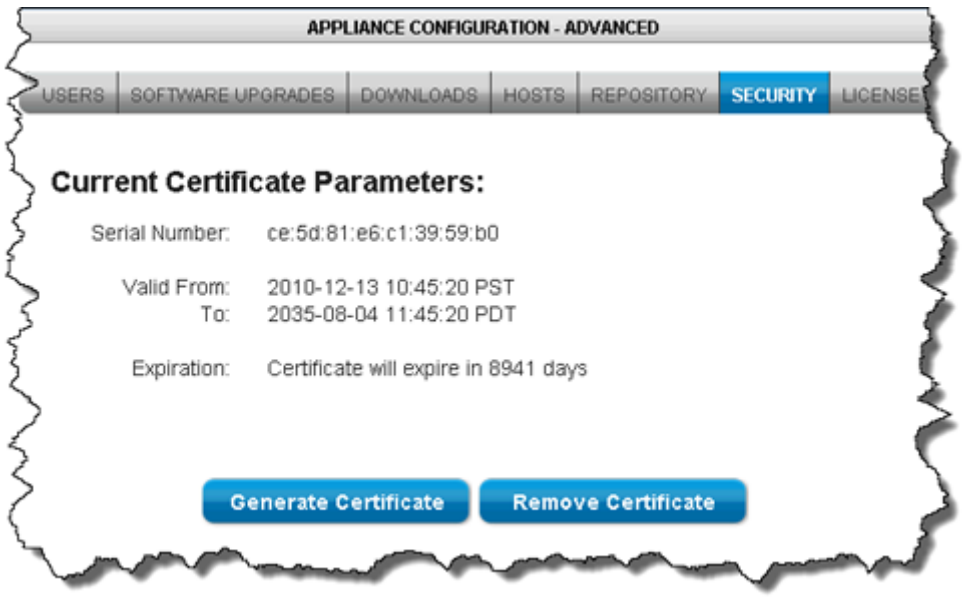
The onQ Appliance uses an X.509 certificate for security. Your onQ Appliance ships with a unique certificate already installed. However, you can create and install a new certificate at any time.

During initial setup of your onQ Appliance, it's a best practice to create a new certificate. Otherwise, there's no reason to create and install a new certificate. The certificate has a long expiration.

To generate a certificate:

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**.

   Don't start protection yet. If you start protection your protected nodes indicate that they have an unsecure connection (see Monitor protected nodes), and this is expected behavior.

5. Install the certificate.

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 them, 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.

9.6 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.

To restart the service:

1. RDP to the Protected Node.
2. From that PN, launch the HA Appliance's onQ Portal.
3. Go to **Start** > **Programs** > **Control Panel** > **Services**.
4. Right-click on **Quorum onQ Service**, and choose **Restart**.

10

Disaster Recovery and DR Preparedness

- Fail over HA Appliance to DR Appliance
- Fail back DR to HA Appliance
- Change Appliance's role
- Start recovery nodes on HA Appliance
- Start recovery nodes on DR Appliance or DR Mirror
- Launch recovery node's console
- Test Mode
- Configure automatic testing of RNs
- Suspend and resume RN tests

10.1 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, Quorum can ship you the cloud DR Appliance, replacing it with a new DR instance. This onsite DR Appliance assumes the HA role. From this Appliance, you can perform the needed BMR/failback to restore your PNs.

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. Add a host entry for each protected node, if you haven't already done so.
5. 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.

Related Topics

Fail back DR to HA Appliance

10.2 Fail back DR to HA Appliance

Use this procedure if you performed the failover steps outlined in 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 two methods:
      - bare metal restore (BMR)
      - 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.

Related Topics

Fail over HA Appliance to DR Appliance

10.3 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 DR Appliance to HA if you experience a failure as outlined in 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.

10.4 Start recovery nodes on HA Appliance

On an HA Appliance there are two ways to start *individual* 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.
- **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 Windows Server updates.

You may be prompted to install application-specific software licenses upon startup of recovery nodes.

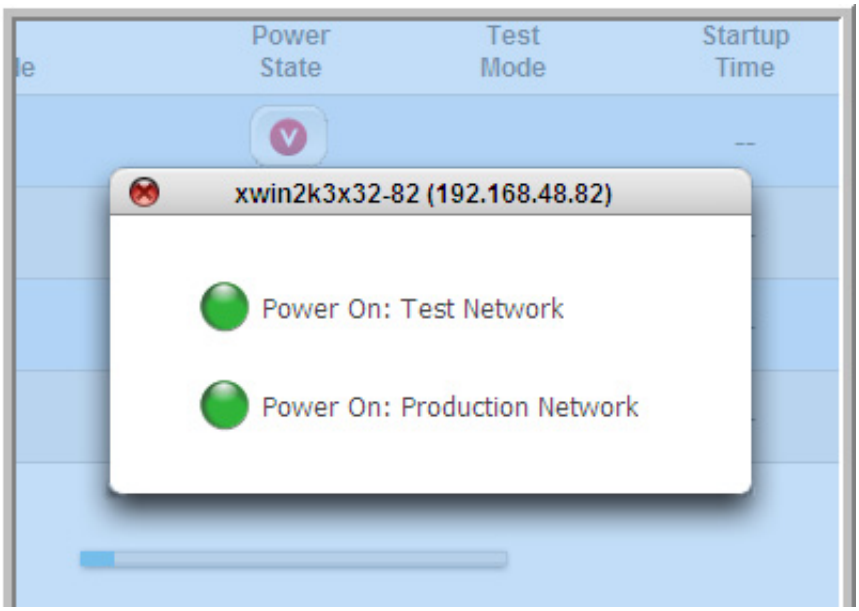
If you want to start *all DR* recovery nodes, go to Start recovery nodes on DR

Appliance or DR Mirror.

(From Recovery Nodes 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. From the HA Appliance, go to **DASHBOARD** tab > **RECOVERY NODES** page.
4. Unlock the page.
5. 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**.
6. Choose the **Power On: Production Network** power option.



7. if you have an Oracle database, you must restore it now. Go to Back up and restore Oracle 11g database.

(From Protected Nodes 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.

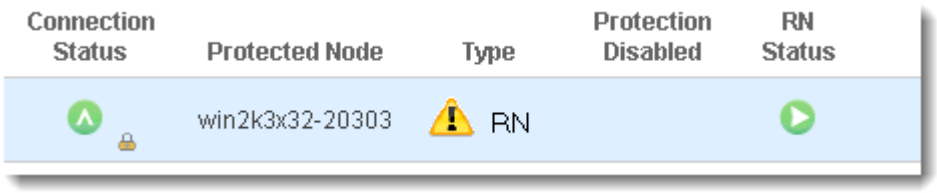
3. Go to **DASHBOARD** tab > **PROTECTED NODES** page.
4. Unlock the page.
5. Click on the**Connection Status** button for the recovery node that you want to start.

   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.

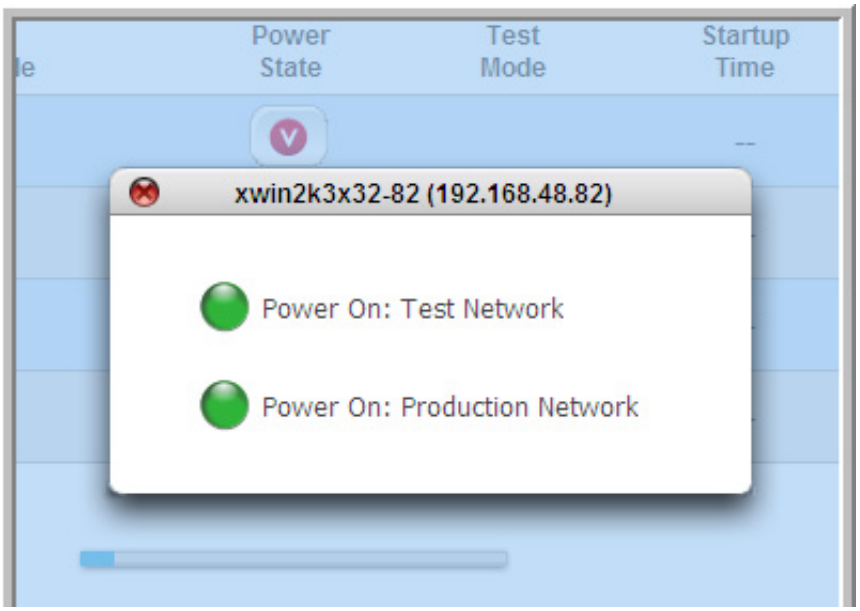
   

7. if you have an Oracle database, you must restore it now. Go to Back up and restore Oracle 11g database.

To start a recovery node in *test mode*:

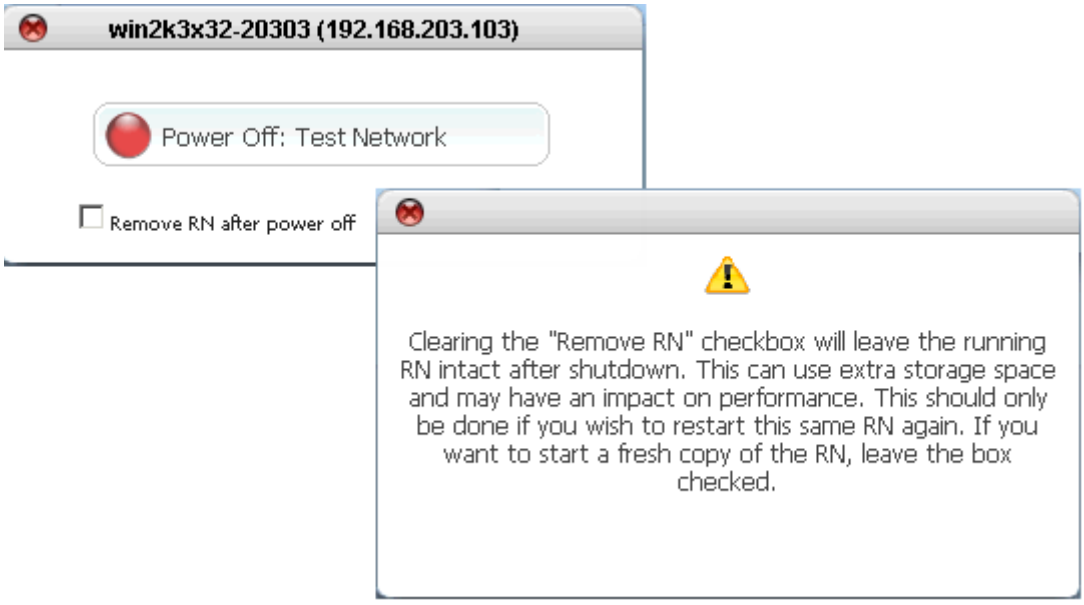
1. Go to **DASHBOARD** tab > **RECOVERY NODES** page.
2. Unlock the page.
3. Locate the **Power State** column.

   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**.
4. Choose the **Power On: Test Network** power option.

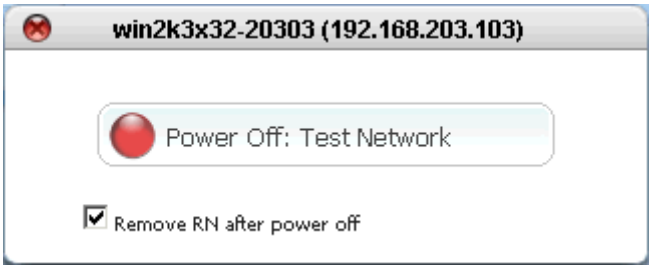


A check mark appears in the **Test Mode** column of the appropriate status page, indicating that the recovery mode is running in test mode.

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.

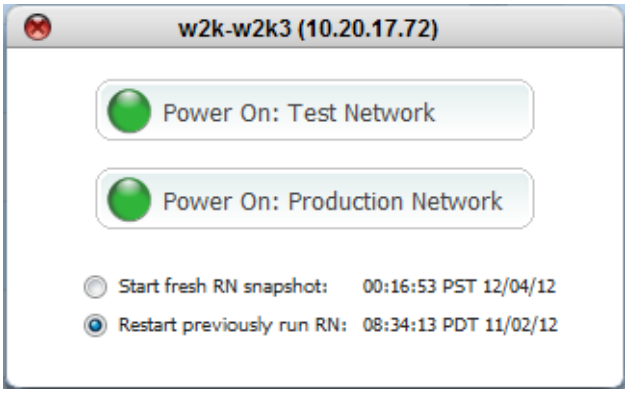


c. Click **Power Off: Test Network** option.



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.

10.5 Start recovery nodes on DR Appliance or DR Mirror

In the event of a disaster, you might need to start all the recovery nodes on the DR Appliance.

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:

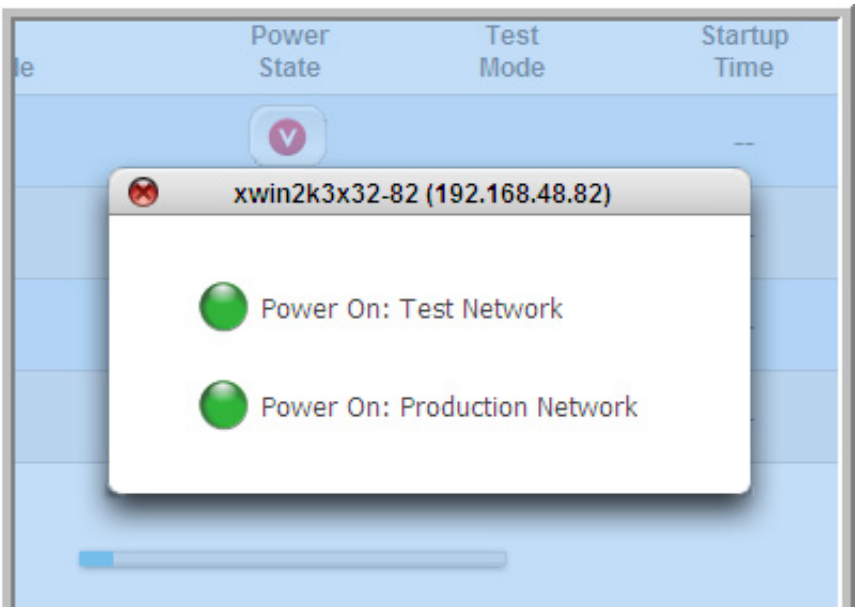
- You can change the memory allocation of PNs so that the total is less than the memory on the onQ Appliance.
- You can start your recovery nodes individually. The boot order must take into account any interdependencies.
- You may be prompted to install application-specific software licenses upon startup of recovery nodes.

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. 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.



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. Locate **Link Status**.
6. Click on the **Link Status** light (and the light is green). You are asked if you wish to **Power On the DR RNs**.
7. Click **Apply** to start all the DR recovery nodes.

10.6 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.
4. Choose **Open with Java Web Start Launcher** to open the.`jnlp` file.

   If your desktop cannot recognize this file, then you don't have JRE (Java Runtime Environment) installed. Download the software from the Java site.
5. If prompted, click **OK** to accept the security warning.

   The **onQ RN Console** window appears.
6. 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

10.7 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 assure, 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

10.8 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. 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

Methods of Performing Incremental Backups

Initiate immediate backups

10.9 Suspend and resume RN tests

RN tests, 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

Monitoring

- Monitor multiple onQ Appliances
- Monitor protected nodes
- Monitor recovery nodes
- Monitor backups
- Monitor DR Appliance
- Monitor onQ Archive Vault enrollment
- Monitor disk space usage
- Interpret e-mail alerts
- Modify e-mail alert settings

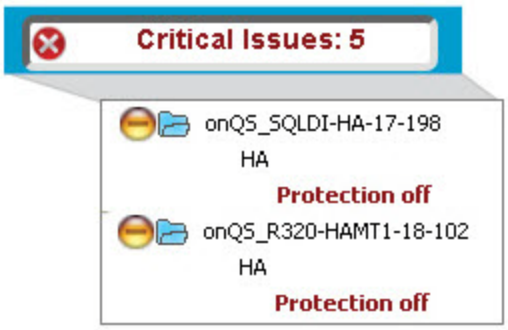
11.1 Monitor multiple onQ Appliances

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 that exist in your environment:
   a. In the navigation pane, click **Add/Remove** button.
   b. Type the address of the HA Appliance, then **Get Name**.
   c. (Optional) Type a **Group** name if you'd like to assign the HA Appliance and its DR, DR Mirror(s), and AV to a group.

      Groups help you organize your appliances. For example, you can organize your appliance by domain, subnet, or department.

      By default, onQ Monitor provides a `quorum` group to include the HA from which you logged on to the onQ Monitor.
   d. Click **ADD**.

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.
3. Wait while onQ Monitor requests status information from the onQ Appliances. 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.
4. From the navigation pane, do one of the following:
   - Select the onQ Appliance to drill down to PNs managed by a specific onQ Appliance.

- Scroll the list to browse status information for all PNs.

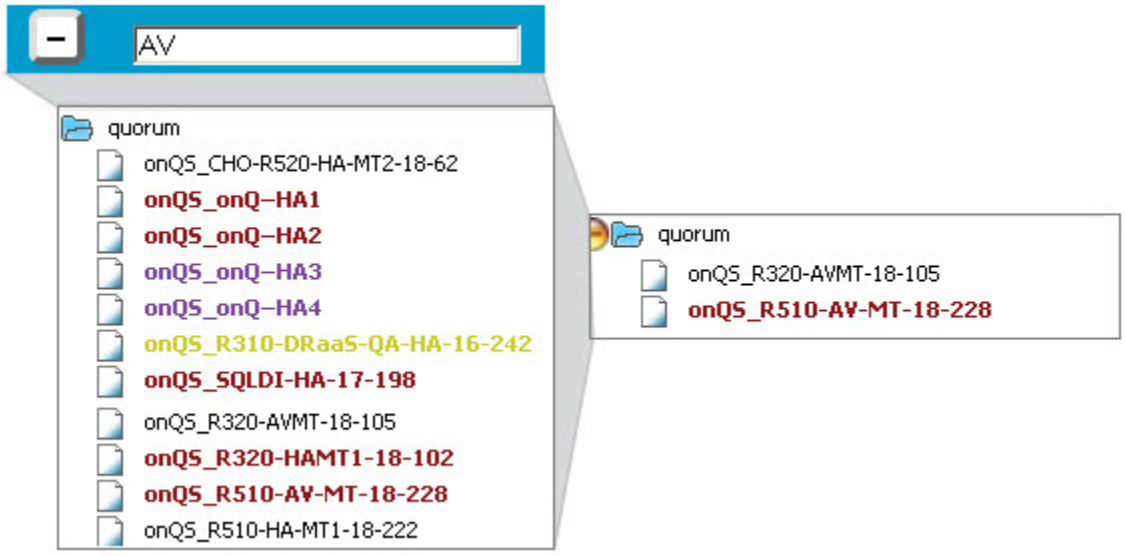
To browse the status of a job:

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.

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.



11.2 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.

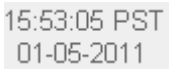
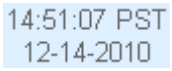
To monitor the status of a protected node:

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 points upward	Protected node is running and protection is turned on.
	 carat is red and points downward	Protected node is not online. Either the protected node is down, or the onQ Service on the protected node is not running (see "Restart the onQ Service" on page 154).
		Appliance's connection to protected node is secure.
		Appliance's connection to protected node is unsecured. See "Troubleshoot Connection Problems" on page 203
Protected Node		Hostname of the protected node. Hostname is the same as the Windows hostname.
Type	**PN**	Protected node itself is running as opposed to the recovery node.
	RN 	Recovery node, 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 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. Hover-over help displays a record of recent events for the RN.
	 carat is green and moves top to bottom	onQ is preparing the Recovery Node
	 carat is red and points downward	No Recovery Node is available either because you just added the PN and onQ is backing up the PN for the first time, Protection is off, or you set the PN to Build on Demand (BoD) by disabling `Auto RN Creation` (see Modify recovery node creation policy). In the latter case, the acronym `BoD` appears next to the carat.
	 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, points upward, and 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. Hover-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 points upward	Complete restorable backup exists in the repository. Hover-over text indicates `Complete`. The date and time of the most recent completed backup appear to the right.
	 carat is green and moves left to right	Backup is in progress. Hover-over text indicates `Running`.
	 carat is green and moves top to bottom	Transfer of the backup to the repository is in progress. Hover-over text indicates `Adding`.
	 carat is green, 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 points downward	Not a single copy of a backup snapshot exists.
	 timestamp	Date and time of the most recent successful backup.
Next Scheduled Backup	 timestamp	Date and time of the next scheduled backup.

Backup Transfer Margin **Note:** First column represents PN backups on HA. Second column represents transfer of those backups to the DR.	 solid green icon	Indicates how well PN backups and DR transfers are keeping up with the backup interval. *Actual time taken was under 65% of the interval* as calculated by averaging intervals from the previous 24 hours (for example, you have a backup interval of 1 hr and your backups completed within less than 39 minutes, on the average); this 24-hr time period automatically resets each time you restart protection. Hover-over help displays: • policy information such as retention period and backup interval • PN's oldest snapshot (in days) • backup transfer times for the last hour.
	 solid yellow icon	Indicates how well PN backups and DR transfers are keeping up with the backup interval. *Actual time taken was 66-89% of the interval* as calculated by averaging intervals from the previous 24 hours; this 24-hr time period automatically resets each time you restart protection. Hover-over help displays: • policy information such as retention period and backup interval • PN's oldest snapshot (in days) • backup transfer times for the last hour
	 solid red icon	Indicates how well PN backups and DR transfers are keeping up with the backup interval. *Actual time taken was 90% or more of the interval* as calculated by averaging intervals from the previous 24 hours; this 24-hr time period automatically resets each time you restart protection. Hover-over help displays: • policy information such as retention period and backup interval • PN's oldest snapshot (in days) • backup transfer times for the last hour.
	 white icon with question mark	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.

Related Topics

Coordination of Backups and Recovery Nodes
Monitor recovery nodes
Monitor backups

11.3 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.
	 carat is red and points downward	Recovery node is not running.
	 carat is yellow and points upward	Recovery node is starting.
	 carat is yellow and points downward	Running recovery node is shutting down
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 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.

Related Topics

Monitor protected nodes

11.4 Monitor backups

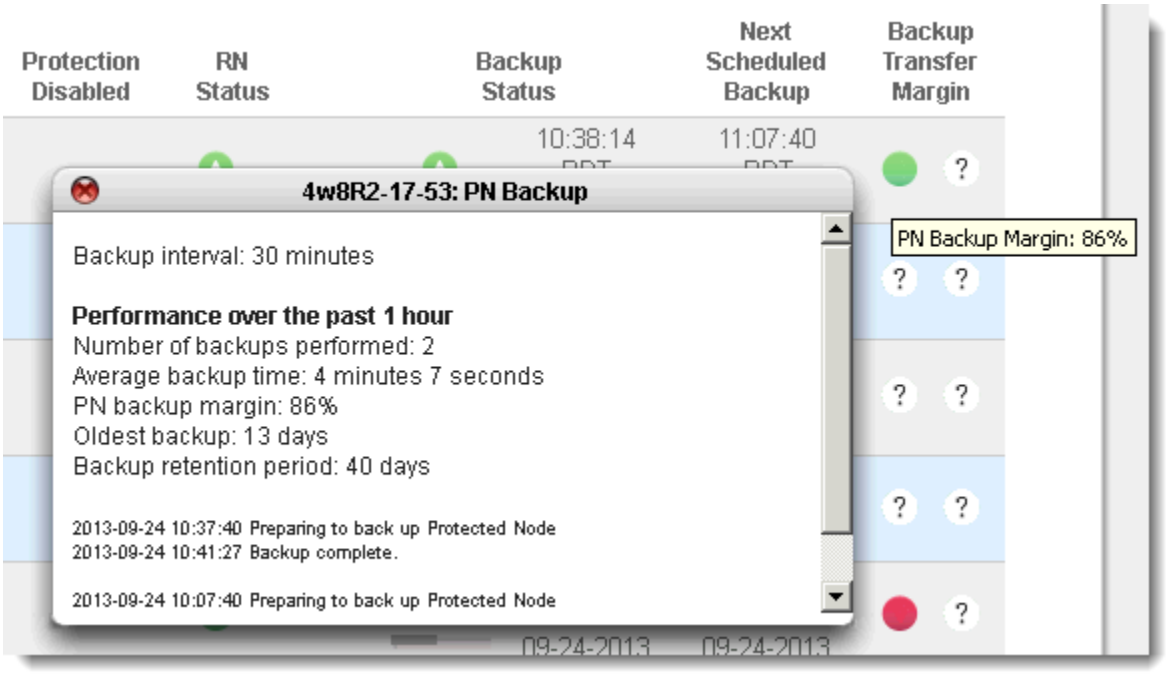
onQ Manager provides two main tools to help you monitor your backups:

- Backup Transfer Margin
- Logs

From a protected node's dashboard, you can get insight into backup transfer margin for that protected node's backups. The backup transfer margin icons indicate how well backups and DR transfers are keeping up with the scheduled interval.

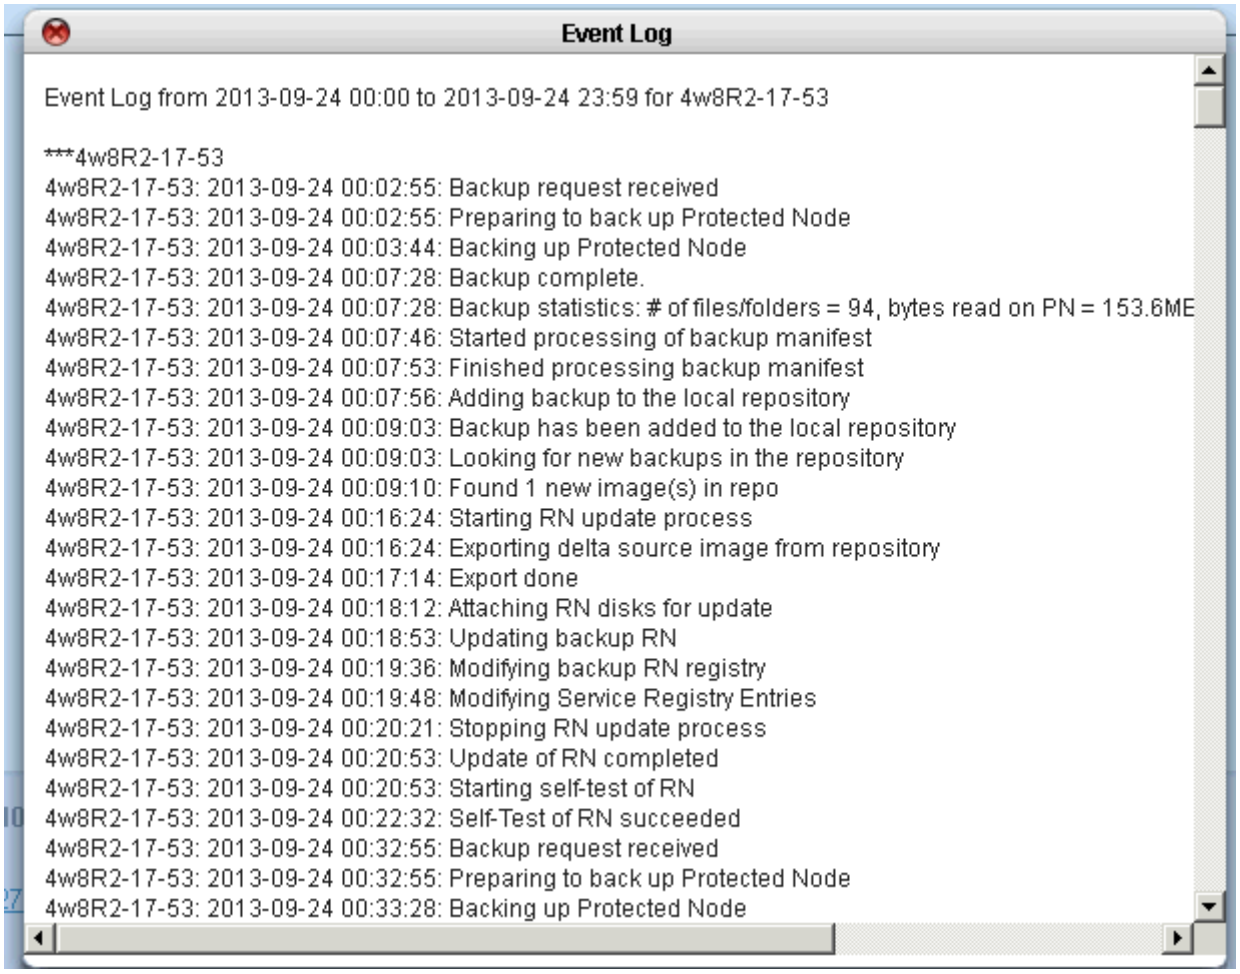
In the event that the interval isn't met, the Event Log indicates `scheduled backup request for <pn_name> occurred before previous backup completed`....In this case, increase your interval time so as to avoid this

scenario.



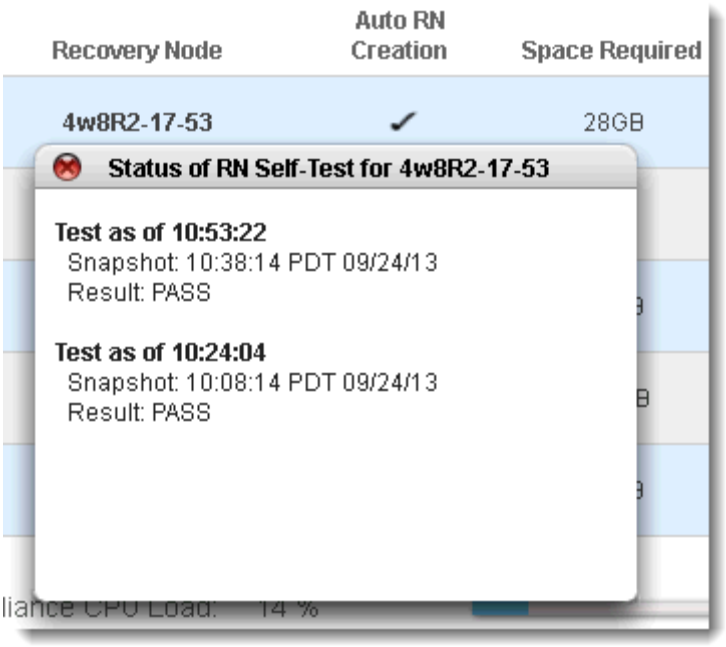
Logs provide a convenient record of internal activity (events) on the Appliance. Logs can serve as a useful means of monitoring and, if

necessary, troubleshooting the Appliance.



There are a few types of logs, but the Event Log is the one you're most likely to browse. If you want to see recent events for individual recovery nodes, simply roll your cursor over the recovery node on either the HA Appliance or

the DR Appliance:



To view a log:

onQ provides several types of logs that you can view or download.

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 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.
- **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 time range for the data you want to display by specifying the start and end dates.
5. 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.

Related Topics

Coordination of Backups and Recovery Nodes

Initial Backup

Backup Activity Report

11.5 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 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.

Hover-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 communication with the DR.
	XFER	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.

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, indicating that the onQ Appliances are not synchronized.
- there's enough disk space to store the snapshots.

Related Topics

Monitor backups

To monitor the DR Appliance:

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.
RN Ready Status	 carat is green and points upward	Recovery node is ready.
	 carat is red and points downward	Recovery node is not ready.
	 hourglass on a yellow background	Preparing to run recovery node.
RN Running Status	 carat is green and points upward	Recovery node is running and protection is turned on.
	 carat is red and points downward	Recovery node is not running and protection is turned off.

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 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).

11.6 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

11.7 Monitor disk space usage

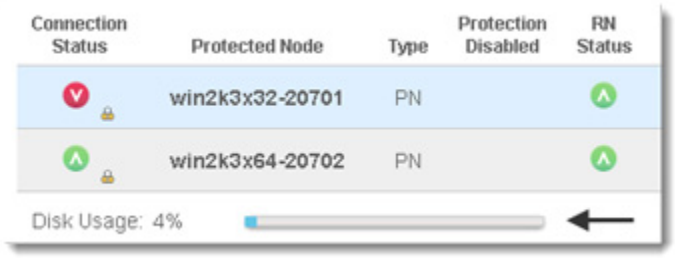
onQ can report on the disk usage of the onQ Appliance Repository and recovery nodes.

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 *recovery 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.

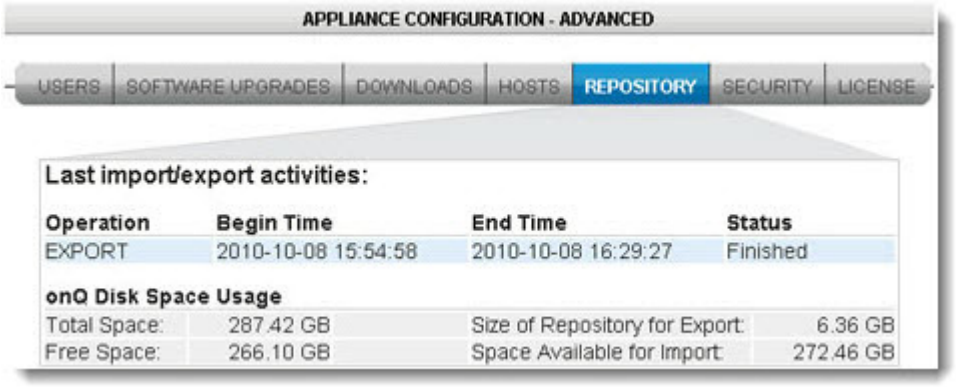


   Disk usage is calculated based on the disk usage of all the protected nodes that are running.

To determine disk space 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

11.8 Interpret e-mail alerts

The onQ Manager sends alerts if it detects possible system failures (for example, if a protected node or an incremental snapshot fails). These alerts can be critical or informational only.

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.

There are several types of alerts:

- License Expiration Alerts
- Reboot Alerts
- Protection Alerts
- vDisk or Disk Space Alerts
- Connection Alerts

11.8.1 License Expiration Alerts

Each Appliance requires a license. Without a valid license, your nodes are not protected. onQ sense license alerts before and after expiration.

Example license expiration alerts are listed in the following table:

Severity & Subject	Message	Cause	What do I do?
NOTIFY – License Expiration Warning	`The installed license will expire on 2010-10-27 11:49:19PDT. You must install a new license before that date in order to maintain protection....`	Threshold met. onQ sends alert everyday 14 to 8 days before expiration	Upload a replacement license. Go to "Install onQ Appliance license" on page 49. ....
CRITICAL – License Expiration Warning	`The installed license will expire on 2010-10-27 11:49:19PDT. You must install a new license before that date in order to maintain protection....`	Threshold met. onQ sends alert everyday 7 to 0 days before expiration	
CRITICAL – License Expiration Alert	`The installed license expired on 2010-10-27 11:49:19PDT. Protection on this onQ appliance has been disabled and will not run until a new license is installed....`	Threshold met. onQ sends alert everyday for 2 days after actual expiration occurs	

11.8.2 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 explanation in the alert. Otherwise, a generic message appears in the alert.

Example reboot alerts are listed in the following table:

Severity & Subject	Message	Cause	What do I do?
NOTIFY – Rebooting onQ	`Process holding the repo lock (3044) cannot be found. The contents of the lock file is: Repo marked dirty on Mon Sep 06 21:24:28 -0700 2010 by 3044 (../src/squirt-add ["--repo", "repo", "--log-level", "INFO", "--ref-file", "/opt/BCVdata/BCVServer/current/addlist.%ALL"])....`	sent once – immediately	Nothing. onQ rebooted the onQ VM and restarted protection. However, If onQ Protection does not start correctly, please contact Quorum Support for assistance.
NOTIFY – Rebooting onQ	`Protection is enabled but the manager process cannot be found....`	sent once – immediately	

11.8.3 Protection Alerts

Attached are the alerts that are sent out when protection is turned on or off on the onQ by a user that has VARAdmin user role.

Severity & Subject	Message	Cause	What do I do?
NOTIFY – Protection has been STARTED on <PN-Name>	onQ protection has been started at user request	A user that has VARAdmin user role started protection on the protected node. onQ sends this alert once – immediately.	Nothing as you must start protection in order to protect your nodes. However, if you did not perform this action, determine who is managing onQ.
NOTIFY – Protection has been STOPPED on <PN-Name>	onQ protection has been stopped at user request	A user that has VARAdmin user role started protection on the protected node. onQ sends this alert once – immediately.	Nothing as there are cases in which you might want to stop protection (see “Stop node protection” on page 78). However, don’t forget to restart protection (see “Start node protection” on page 78). Also, if you did not perform this action, determine who is managing onQ.

11.8.4 vDisk or Disk Space Alerts

A protected node's virtual disk (vdisk) requires a specific amount of disk space. As soon as your protected node's vdisk capacity usage exceeds 80%, onQ alerts you. Also, your onQ Appliance has internal policies that ensure that it doesn't run out of disk space.

Severity & Subject	Message	Cause	What do I do?
CRITICAL – Failure of Recovery Node Creation for <PN-Name>	`An unexpected error was encountered when attempting to update the contents of a Recovery Node. The error status returned was: SystemErrDiskFull(e:). No further updates to the Recovery Node will be attempted until you select the Initialize button. Please contact Quorum Support to resolve this error.`	One of two reasons: • RN was built for the first time and vdisk for one or more volumes is under allocated. • RN was re-initialized (even if earlier RN exists) and vdisk for one or more volumes is under allocated.	Resize/expand the vdisk as outlined in "Resize protected node's vdisk" on page 87. The error message indicates the relevant disk drive/volume. In this example, `e:`.

Severity & Subject	Message	Cause	What do I do?
CRITICAL – RN disk space insufficient for <PN-Name> on <onQ-Manager-Name>	The Recovery Node disk space allocated for <PN-Name> is insufficient to create a new RN from the current snapshot. The previous RN will be left unchanged. The snapshot data for the PN is intact and a new RN will be created when the configured disk space is increased....	Initial vdisk size allocated for a volume on the PN (when PN was configured via the Protect Me button or manually on to the HA) is inadequate because of the growing data on PN in the corresponding volume. In this case, the onQ Manager does not update the RN. You will continue to receive this alert until you take corrective action.	Resize the vdisk as outlined in "Resize protected node's vdisk" on page 87.
CRITICAL - <onQ Appliance-Name>: Free Diskspace Was Critically Low	The free Repository space on onQ-HASSD has fallen below allowable limits. OnQ expired oldest 10% of backups (total of 8) for each Protected Node in order to recover space: <PN-Name-16=>7 <PN-Name-93=>1	onQ has a built-in policy whereby if its disk space exceeds 85% onQ begins to "free up" space by deleting old backups.	Add more disk space, or delete backups manually as outlined in Delete snapshots,

Related Topics

Backup Activity Report

11.8.5 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.

Severity & Subject	Message	Cause	What do I do?
Critical CRITICAL - Link between <HA-Appliance-Name> and <DR-Appliance-Name> is re-established	`The link between the HA and DR sites has been re-established. If RNs were being used in production mode, you will need to schedule your failback procedure before bringing up your PNs to minimize data loss.`	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.	If necessary, go to "Fail back DR to HA Appliance" on page 158
Critical CRITICAL - Unable to contact <HA-Appliance-Name> from <DR-Appliance-Name>	`The onQ HA Manager onQMspHA7.cust7.com cannot be contacted. If this is confirmed to be a site failure, you will need to start Disaster Recovery protection from MspDR7.`		If necessary, go to "Fail back DR to HA Appliance" on page 158

11.9 Modify e-mail alert settings

onQ sends alerts if it detects possible system failures. To receive such alerts, you must configure the 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' 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 the email addresses for the recipients of the two different alert severities: `Notify` and `Critical`.

   Type multiple email addresses on the same line, separated by commas or spaces.

6. Click **Send** 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

Troubleshooting

- Troubleshoot Installation Problems
- Troubleshoot Connection Problems

12.1 Troubleshoot Installation Problems

Use the following table to help you troubleshoot installation problems.

Symptom	What's it mean?	What do I do?
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, which is okay.	If you intend for them to be on different subnets, ignore this message. This message is just a reminder.
I don't see the **Protect Me** button.	The following four conditions must be met before the **Protect Me** button is visible: • The onQ Appliance must have been initially configured. • The node's OS must be supported by onQ. • Protection must be turned **OFF**.	Verify that the conditions exist.

12.2 Troubleshoot Connection Problems

Use the following table to troubleshoot connection problems.

Symptom	What's it mean?	What do I do?
PN's Connection Status icon is red	• IP Address Change • Host name change	Update IP address and hostname. See Manage hosts.
	DNS Server is down	Bring up DNS Server. Add host. See Manage hosts.
	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.
	Certificate is corrupt	Regenerate and install the certificate. See Secure HA connection to PNs.
	onQ Service is not running	Restart the service. Go to Restart the onQ Service.

13

Windows Cluster Services

- Windows Cluster Services Overview
- Download iSCSI initiator for Windows 2003
- Configure Quorum onQ Service
- Work with RNs in Windows Cluster Services Environment

13.1 Windows Cluster Services Overview

onQ supports the ability to run recovery nodes in a Windows Cluster Services environment. How you configure, start, and run the cluster services on a Recovery Node is different than in a non-cluster environment.

For additional information on cluster support, refer to the **Enable Cluster Support** parameter outlined in Protected Node Parameters.

Related Topics

Windows Cluster Services Limitations

13.2 Download iSCSI initiator for 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.

To download the initiator:

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`

13.3 Configure Quorum onQ Service

To configure the Quorum onQ Service (*Windows 2012*):

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**.

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

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**.

To configure the Quorum onQ Service (*Windows 2003*):

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**.

13.4 Work with RNs in Windows Cluster Services Environment

Before performing these procedures, make sure to complete the Configure Quorum onQ Service procedure.

To configure the onQ Appliance:

1. Turn onQ Protection OFF.
2. Upgrade onQ to 3.6 HF1, if not already at that release level.
3. Putty or Terminal into the hypervisor
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 XenCenter'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.

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**.

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, re-initialize the RN to trigger a complete RN build.

   If the PN is 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 2008 R2* & 2012):

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 160 or "Start recovery nodes on DR Appliance or DR Mirror" on page 165.
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):
   - To start the RN in test mode (Windows 2008):
   - To start the RN in test mode (Windows 2008 R2 & 2012):

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 audit reasons, this 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.

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 *Admin, 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

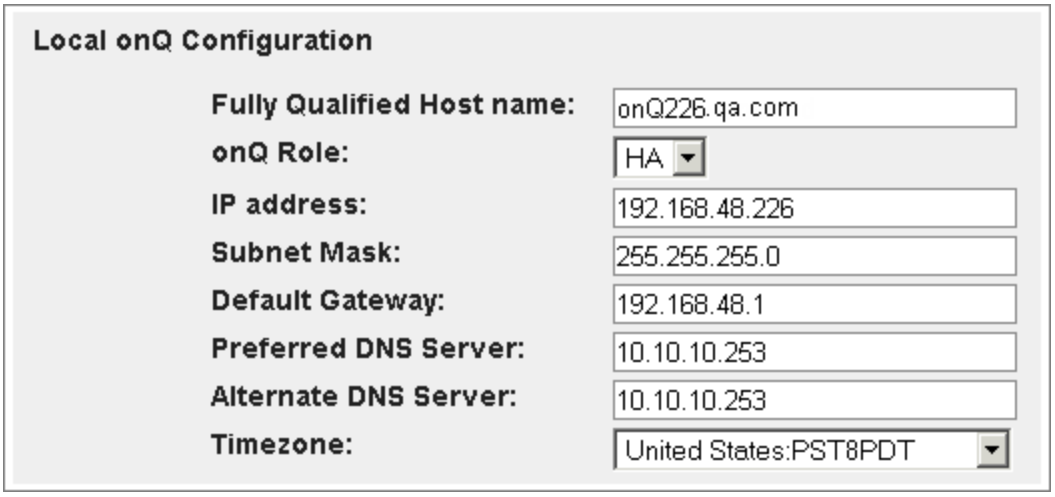
Example Configurations

- 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

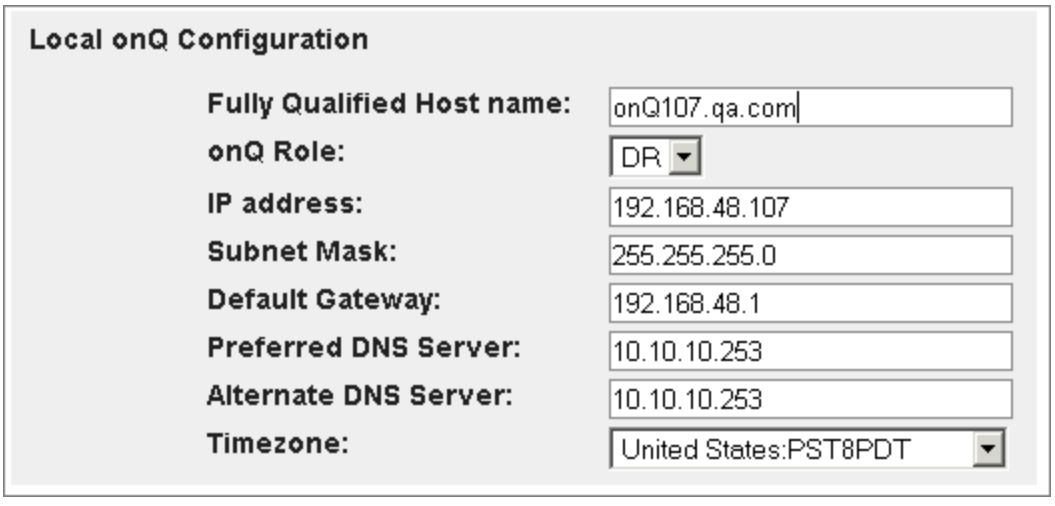
15.1 Local onQ Appliance Configurations

- Local: Example of DR Appliance
- Local: Example of DR Mirror

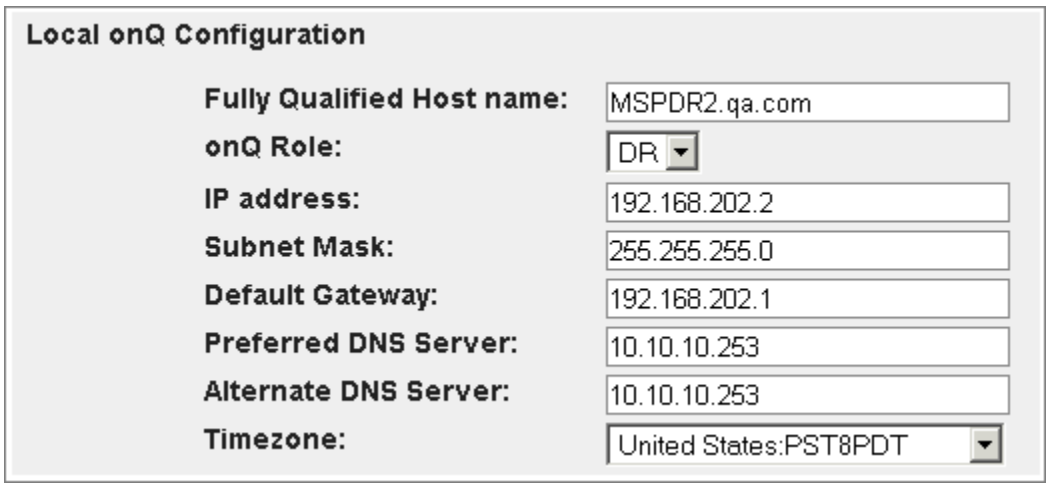
15.1.1 Local: Example of HA Appliance



15.1.2 Local: Example of DR Appliance



15.1.3 Local: Example of DR Mirror



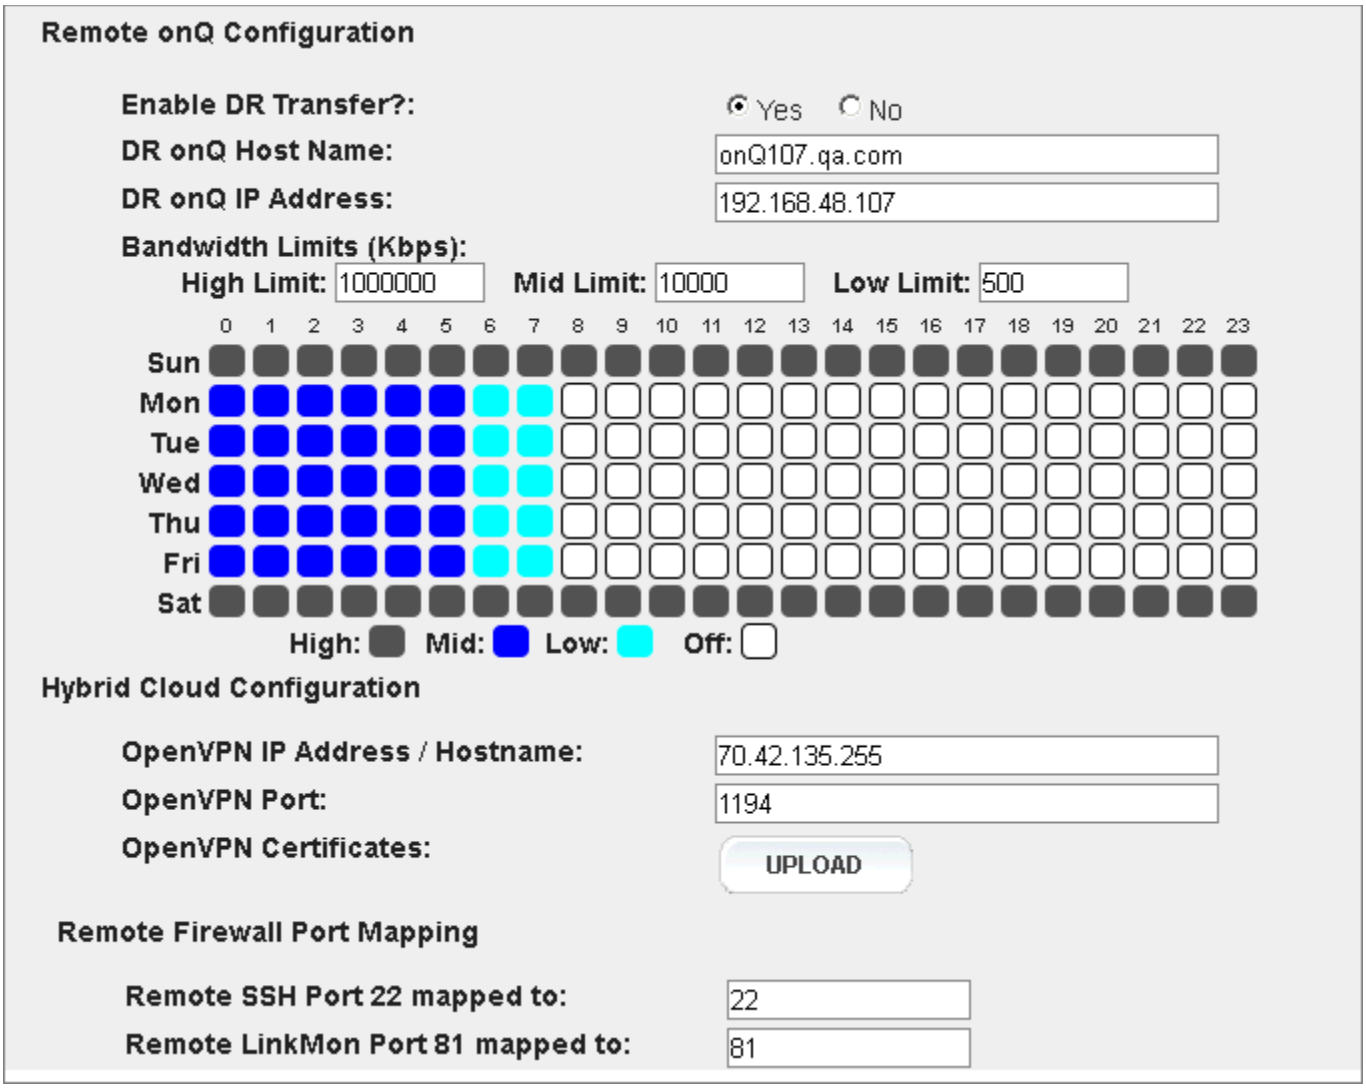
15.2 Remote onQ Appliance Configurations

- 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

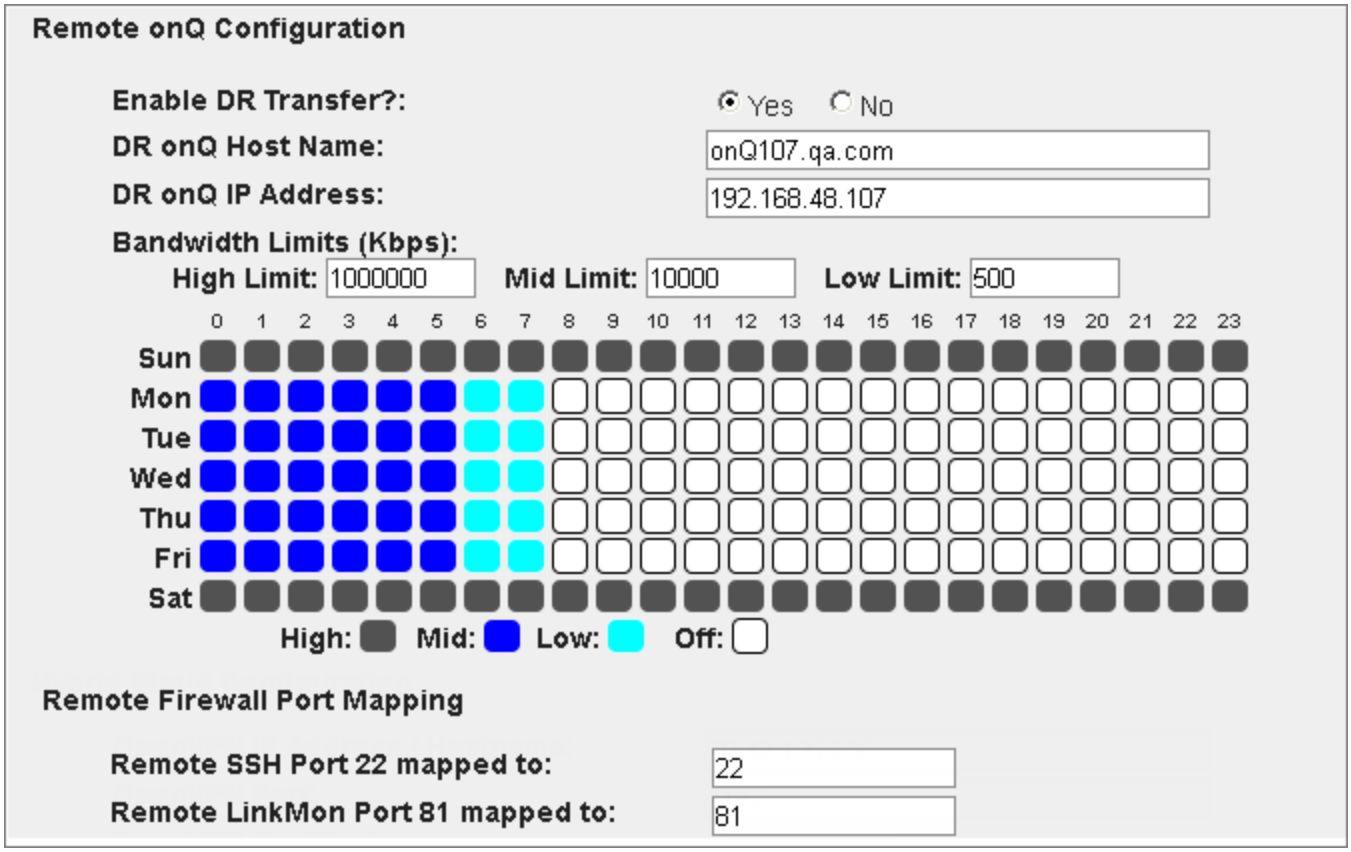
In order for the HA Appliance to communicate with the DR Appliance, via a secure VPN connection, you must provide the following information for the remote Appliance.

- 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

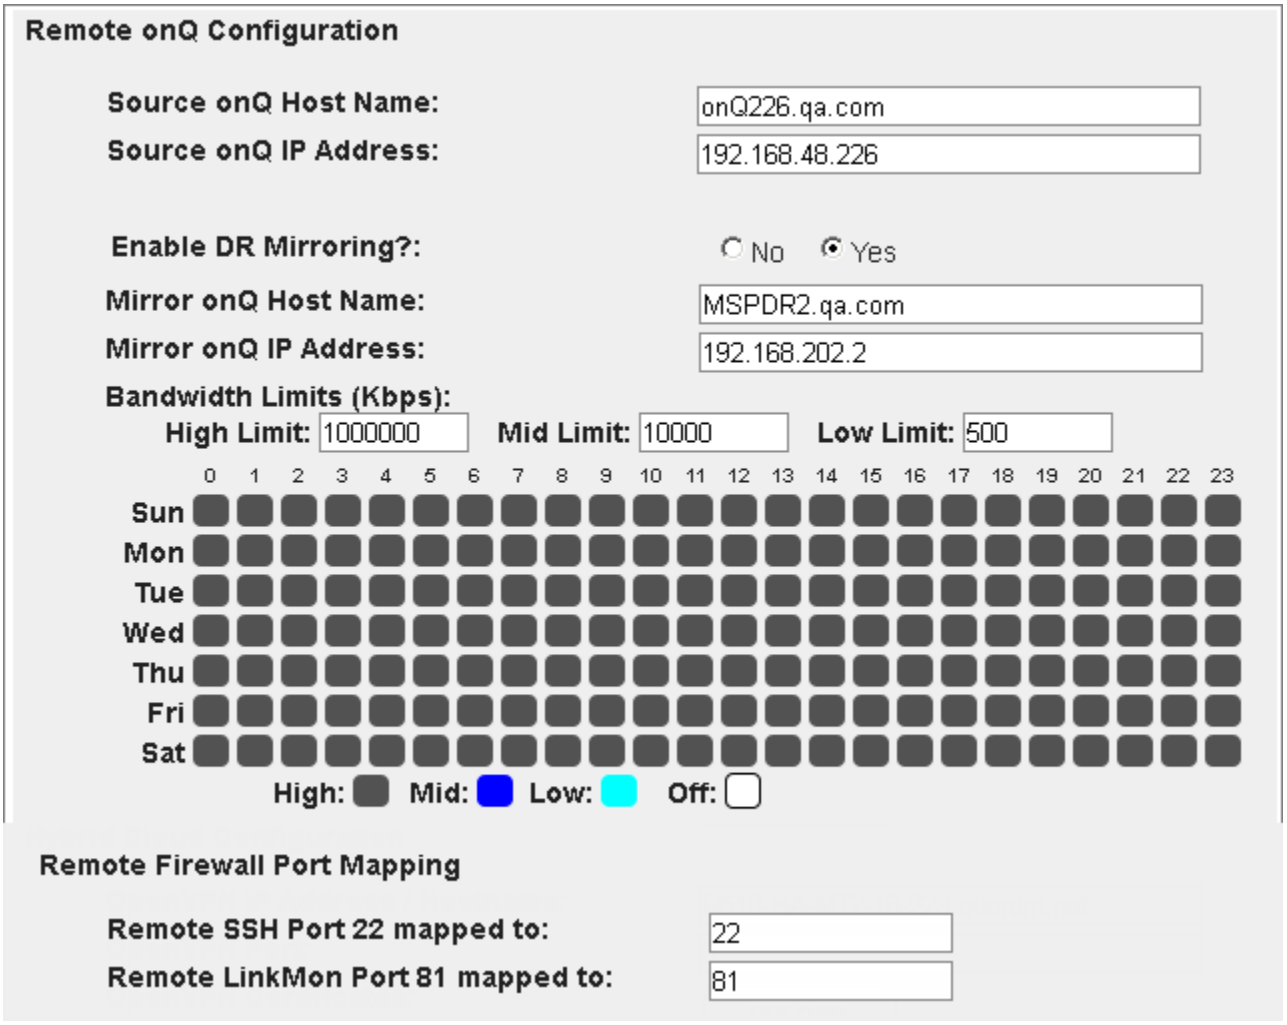
Warning: Enable DR Transfer?. After an initial configuration, do not disable DR transfers as outlined in Disable snapshot transfers. If you do so, your repositories will be out of sync.



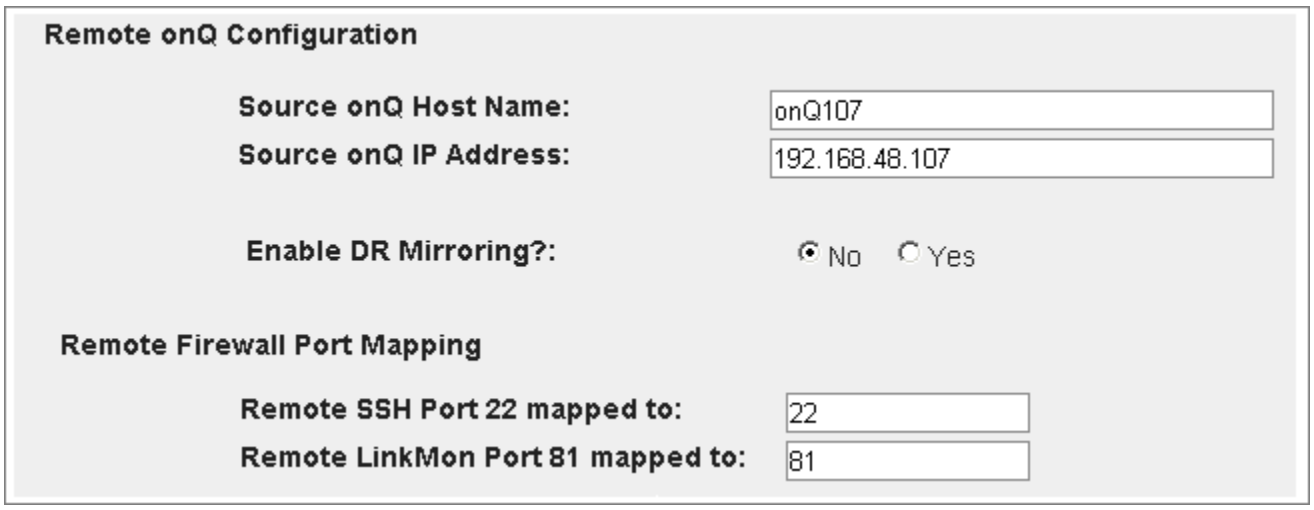
15.2.3 Remote: Example of DR Appliance with DR Mirror

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

Tech Notes

- Search for EFI Partition
- Make EFI Partition Bootable
- Install kernel-xen RPM package

16.1 Search for EFI Partition

You probably are already familiar with your system's hardware. However, if you're unsure if you have an EFI partition, perform this procedure. QUARK's user interface does not support EFI partitions. If you have an EFI partition, part of a BMR (Bare Metal Restore) using QUARK must be performed manually.

To determine if you have an EFI partition:

1. Log on to the PN.
2. Run the following command:

```
# bcdedit
```

3. Compare the output with the following output. If the disk has an EFI partition, the path for the disk has `.efi` extension.

GPT = **Yes**	C:\>bcdedit Windows Boot Manager -------------------- identifier {bootmgr} device boot description onQ Windows default {current} displayorder {current} timeout 30 Windows Boot Loader ------------------- identifier {current} device partition=C: path \Windows\System32\Boot\winload.efi description Windows server 2008 R2 osdevice partition=C: systemroot \Windows resumeobject {b82bb1a7-8a52-11e1-9bc6-806e6f6e6963} detecthal Yes

GPT = **No**	C:\>bcdedit Windows Boot Manager -------------------- identifier {bootmgr} device partition=\Device\HarddiskVolume1 description Windows Boot Manager locale en-US inherit {globalsettings} default {current} resumeobject {6d4ead60-6732-11e1-9b71-c3428df8f48b} displayorder {current} toolsdisplayorder {memdiag} timeout 30 Windows Boot Loader ------------------- identifier {current} device partition=C: path \Windows\system32\winload.exe description Windows Server 2008 R2 locale en-US inherit {bootloadersettings} recoverysequence {6d4ead62-6732-11e1-9b71-c3428df8f48b} recoveryenabled Yes osdevice partition=C: systemroot \Windows resumeobject {6d4ead60-6732-11e1-9b71-c3428df8f48b} nx OptOut

4. Return to Perform Bare Metal Restore. That procedure walks you through both scenarios.

16.2 Make EFI Partition Bootable

QUARK's user interface does not support EFI partitions. If you have an EFI partition, part of a BMR (Bare Metal Restore) using QUARK must be performed manually.

To make EFI partition bootable:

1. Check the size of EFI system partition (ESP) and Microsoft reserved partition (MSR), and record this information from your original PN

```
> diskpart
> sel disk 0
> list part
```

Note: Optionally, you can retrieve the same information from the saved image size (ESP.dat).

Example:

```
Partition 1, system partition, 100M

Partition 2, Reserved partition, 128M

Partition 3, Primary partition, 97G
```

2. Back up PN to onQ repository. Refer to "Initiate immediate backups" on page 105.

   The PN's EFI partition (current) will be saved as `C:\Program Files\Quorum\usr\ESP.dat`.

3. Prepare the restore target partitions on GPT disk:

   a. If you are using the same hardware but replacing the drives, ensure that the original RAID disk is removed and recreated with a new set of disks.

   b. Reboot the machine and set the system to use BIOS-mode boot.

   c. Change to BIOS-mode boot.

   d. Boot the restore target with the QUARK USB drive.

Do the following:

```
> diskpart
> sel disk 0
> clean
> covert GPT
```

e. Remove any automatically created partition, if any.

```
> sel part 1
> del part override
```

f. Create an EFI system partition (ESP), MSR partition, and C volume in sequence.

The volume sizes of ESP and MSR must match the sizes on the original PN according to Microsoft's documentation. The size of C volume can be any size that is able to accommodate the restored files on the C volume.

```
> sel disk 0
> create part efi size=100
```

The newly created ESP size must be equal to or larger than saved ESP image (ESP.dat).

```
> format fs=fat32 quick
> create part msr size=128
```

g. Use QUARK to add C volume (required size) without checking the `active` option.

4. Perform BMR on all protected volume as usual. After the data transfer is complete as indicated by onQ, perform the rest of the steps below before rebooting the newly restored server.

5. Restore the saved EFI partition from the C volume to the EFI partition on RN's GPT disk.

```
> diskpart.exe
> sel disk 0
> sel partition 1
```

a. Assign the EFI partition to an available disk letter.

```
> assign letter=p
```

b. Restore the saved ESP image.

```
> B:\BCV\qRestorePart.exe c:\Program Files\Quorum\usr\ESP.dat P:
```

6. Fix BCD store to point to new device path:

a. Use `Bcdedit` to browse the ESP's BCD store, then fix the `device/osdevice` path.

b. Look for the device or osdevice parameter for unknown or a missing path parameter.

```
> Bcdedit /store p:\EFI\Microsoft\Boot\BCD /enum all
```

c. Set the device to the drive letter for the EFI System Partition (ESP).

```
> Bcdedit /store P:\EFI\Microsoft\Boot\BCD\ /set {bootmgr} device partition=p:
```

d. Ensure that the Boot Loader device and osdevice settings point to the correct drive letter for the operating system partition.
   - Where the operating system device letter is usually `C:\`.
   - Where `ID GUID` is the Windows Boot loader identifier. See output from Step b.

```
> Bcdedit /store p:\EFI\Microsoft\Boot\BCD /set {<ID GUID>} device partition=C:
> Bcdedit /store p:\EFI\Microsoft\Boot\BCD /set {<ID GUID>} osdevice partition=C:
```

e. (Optional) Check for any other unknown device path. Update the value to partition=p: or partition=c, according to its correctly assigned driver letter. The same fix can be applied to the multi-boot BCD set up as well.

7. Reboot machine and hit F2 to switch back to EFI-boot mode. Use `diskpat` to verify the system partition is pointed to ESP.

16.3 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
```

17

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 a 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.

squirtcopy ..... onQ's backup utility.

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.