Quorum®

onQ Release Notes

Version 3.8 SP1

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

Copyright © 2008–2015 QuorumLabs, Inc.

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

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

TRADEMARKS

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

Table of Contents

Preface 1

Online Help 1

Chapter Organization 1

1. What's New 3

2. Support 7

2.1 Platform Support 7

2.2 Hypervisor Support 8

2.3 Browser Support 9

2.4 Bare Metal Restore Support 9

2.5 Ontrack Support 10

2.6 Windows Cluster Services Support 10

2.7 Archive Vault Configuration Support 10

2.7.1 Standalone Configuration 12

2.7.2 DR-only Configuration 13

2.7.3 Redundant Configuration 14

2.7.4 Shared-Instance Configuration 15

2.7.5 Unsupported Configurations 15

3. Limitations 17

3.1 User Interface Limitations 17

3.2 Linux RN Limitations 19

3.3 Windows Server Limitations 20

3.4 Volume Mount Limitations 21

3.5 CIFS Mount Limitations on Windows 21

3.6 Bare Metal Restore Limitations 22

3.7 File-level Restore Limitations 23

3.8 Windows Cluster Services Limitations 23

3.9 Agent-less PN Enrollment Limitations 24

4. Updates 27

4.1 Upgrade Support 27

4.2 Update Appliance software 28

Preface

This guide is intended for System Administrators who need to keep abreast of new onQ features and ongoing support and who want to upgrade their onQ Appliance.

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/3.8SP1/onQ/index.html
AV	http://www.quorum.net/help/3.8SP1/Vault/index.html

Chapter Organization

The chapters and appendices of this manual are as follows:

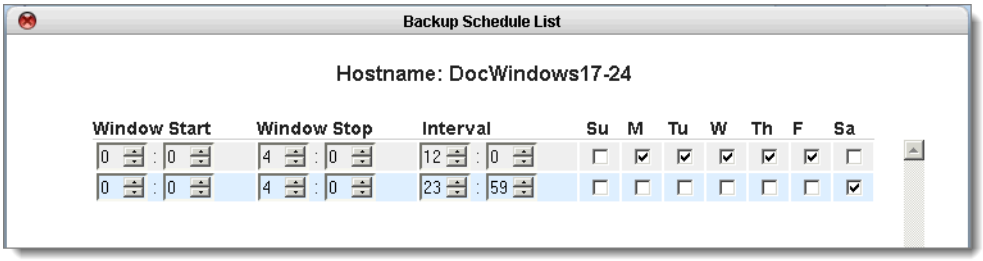
- Chapter 1, *What's New*
- Chapter 2, *Support*
- Chapter 3, *Limitations*
- Chapter 4, *Updates*

1

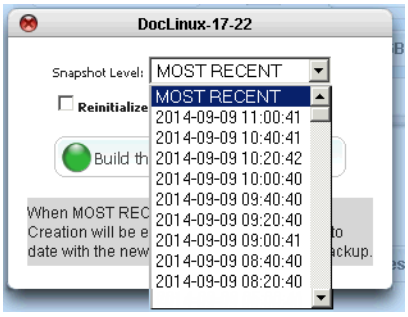
What's New

This release includes the following new features and enhancements:

- **Ability to disable replication**. You can now disable replication for protected nodes that are not important for business continuity. In doing so, onQ will still generate backups for them and archive the backups as instructed. For more information, go to "Protected Node Parameters" in onQ Administrator's Guide.
- **Granular backup scheduling**. You now have more control than ever over your backups. If your network is plagued with low bandwidth, take advantage of granular scheduling. For more information, go to "Schedule backups" in onQ Administrator's Guide.

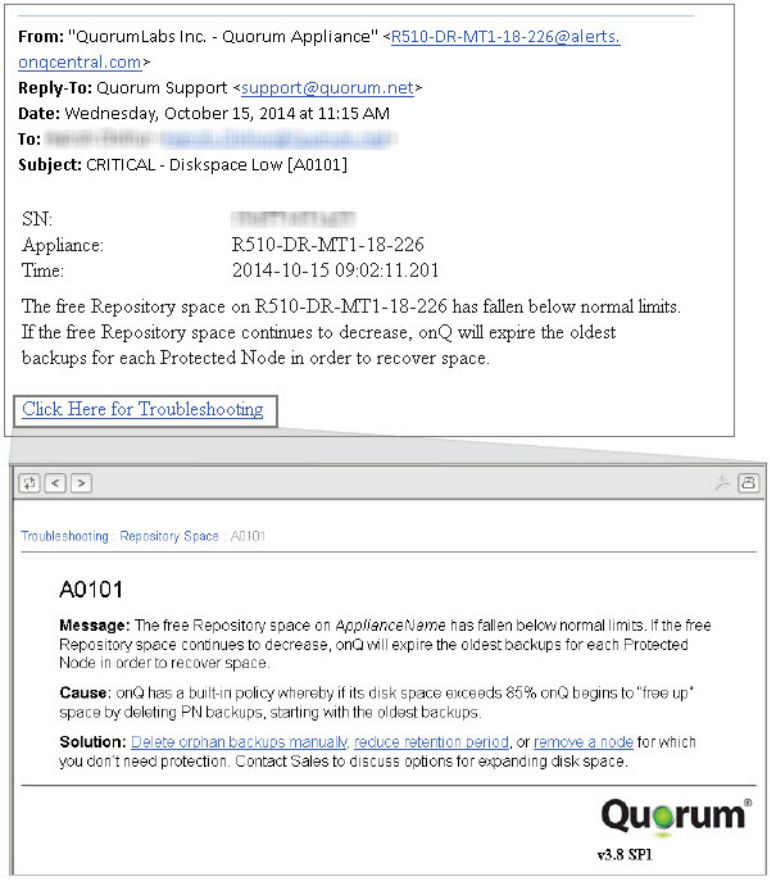


- **Snapshot selection from RN builds**. You can now choose to build an RN from a given snapshot. You are no longer restricted to building an RN from the most recent snapshot. For more information, go to Recreate recovery nodes.



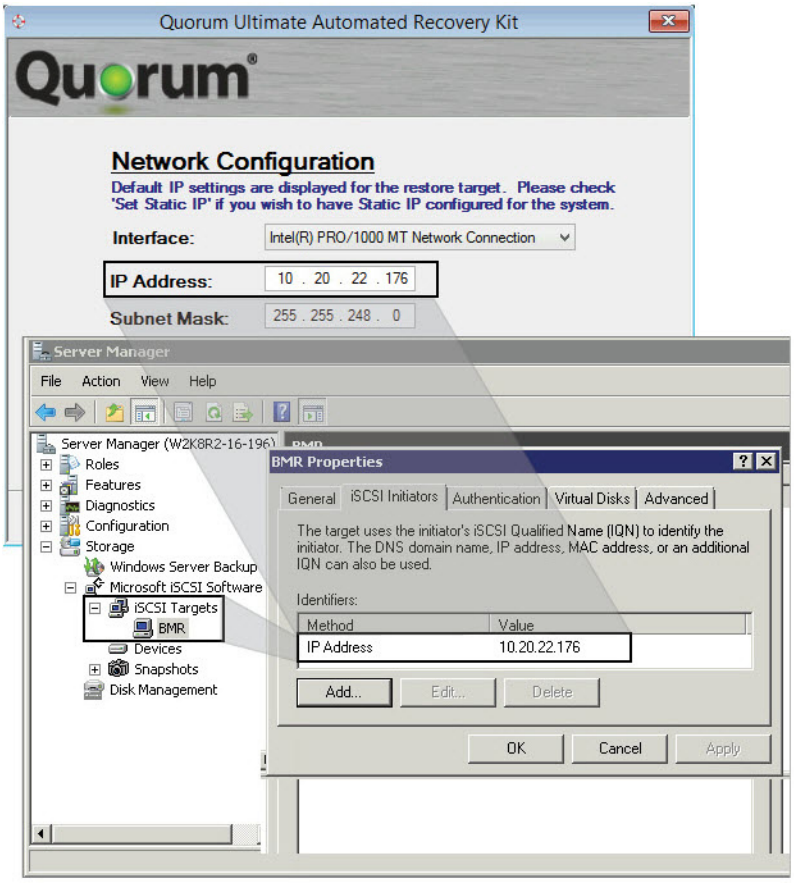
- **Context-sensitive Troubleshooting**. Email alerts now provide you context-sensitive troubleshooting tips. Simply click on the Click Here for

Troubleshooting link in the email alert to be directed to the related Online Help topic. Where applicable, you'll be given tips to help resolve issues with your PNs or onQ Appliance, including backup failures and disk space errors.



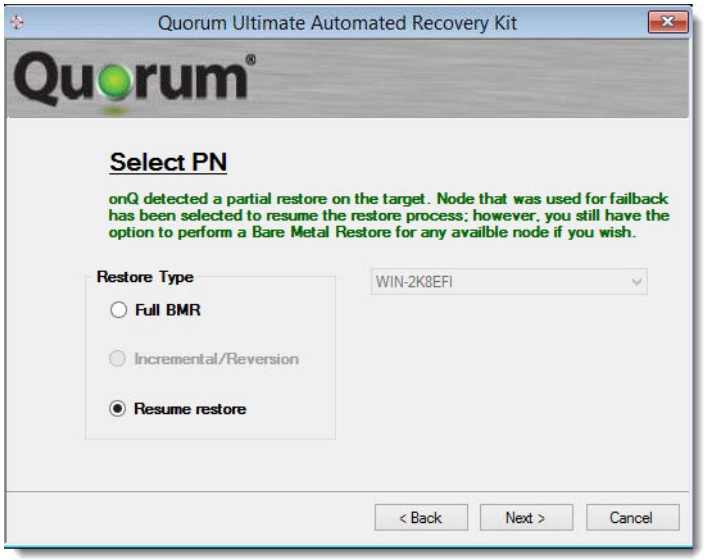
- **onQ Flex**. A new configuration that provides you Instant Recovery for those system that are critical to business continuity and Recovery on Demand for all other systems. This configuration has a couple nuances that you need to know about: see About onQ Flex.
- **iSCSI support**. onQ can now protect and restore iSCSI devices. In the case of a restore, simply add QUARK (BMR target) to the iSCSI target's initiator list, add the iSCSI devices to QUARK's Disk Configuration page,

then leave the rest up to QUARK!. For more information, go to "(Step 2) To restore data to your PN:" on page 275.



- **Ability to resume a restore**. In the event that a restore fails due to a 48-hr timeout issue, or any other reason, the onQ detects this restore failure. In turn, the wizard enables you to resume the restore, "beginning where

QUARK left off". For more information, go to "(Step 2) To restore data to your PN:" on page 275.



2

Support

- Platform Support
- Hypervisor Support
- Browser Support
- Bare Metal Restore Support
- Ontrack Support
- Windows Cluster Services Support
- Archive Vault Configuration Support

2.1 Platform Support

Both physical and virtual PNs are supported on the following platforms:

Operating System[a]	Backup Supported	RN Supported
Windows		
Windows Server 2012 R2, Datacenter and Standard, 64-bit	Yes	Yes
Windows Server 2012, Datacenter and Standard, 64-bit	Yes	Yes
Windows 2008 R2, 64-bit	Yes	Yes
Windows 2008, 32-bit and 64-bit	Yes	Yes
Windows 2003 R2, 32-bit and 64-bit	Yes	Yes

Operating System[a]	Backup Supported	RN Supported
Windows 7, 64-bit	Yes	Yes
Linux		
Red Hat Enterprise Linux/CentOS[b] 5.6, 5.7, 5.8, 5.9, 6.0, 6.1, 6.2, 6.3, 6.4; 32-bit and 64-bit	Yes	Yes
Oracle Enterprise Linux 5.8, 32-bit and 64-bit	Yes	Yes

a.PNs running on OEM licenses can be backed up; however, the RNs may fail to fully initialize because OEM licenses are not transferable to virtual machines.
b.Quorum tested onQ with CentOS 6.4, 32-bit and 64-bit. We expect all other versions to work too.

Related Topics

Limitations
Hypervisor Support
Browser Support

2.2 Hypervisor Support

The guest operating systems outlined in "Platform Support" on page 7 are supported on the following hypervisors:

Hypervisor	Version
VMWare ESX	4.0
VMWare ESXi	5.0 and 5.1.0
Hyper-V	6.x
XenServer	5.0, 5.5, and 6.0.2

2.3 Browser Support

onQ supports the following browsers; however, there are some limitations:

Browser	onQ Portal	Agent-based PN Enrollment (“Protect Me” button)
Mozilla Firefox 29.0 and above	Yes	Yes
Mozilla Firefox 10.0-28.0	Yes	Yes (32-bit) No (64-bit)
Internet Explorer 10 and 11	Yes	Yes
Internet Explorer 7, 8, and 9	No	Yes

2.4 Bare Metal Restore Support

A Bare Metal Restore (BMR) is supported on the following platforms:

Operating System	To Same Hardware	To Dissimilar Hardware
Windows Server 2012 R2, Datacenter and Standard, 64-bit	Yes	Yes
Windows Server 2012, Datacenter and Standard, 64-bit	Yes	Yes
Windows 2008 R2, 64-bit	Yes	Yes
Windows 2008, 32-bit and 64-bit	Yes	Yes
Windows 2003 R2, 32-bit and 64-bit	Yes	Yes

2.5 Ontrack Support

As outlined in Perform WSR Restore, you can use Kroll's Ontrack application to restore Exchange Server data and mailboxes.

Ontrack is independent of onQ; as such, onQ should work with any version of Ontrack so long as onQ supports the underlying operating system (see Platform Support).

Kroll OnTrack Power Controls for Exchange application version 7.0.x is required for using Ontrack with Exchange 2013, in order to support EDB files.

2.6 Windows Cluster Services Support

MSFT for Windows 2012, Windows 2008 R2, Windows 2008, and Windows 2003 are supported. onQ supports iSCSI connections to a storage device in a SAN environment.

onQ supports an active/passive arrangement, which is sometimes referred to as a Single Copy Cluster (SCC). SCC is one of four high availability options of Exchange 2007. Exchange 2010 does not support SCC. With Exchange, only the mailbox role is supported in a cluster environment. Therefore, cluster support for Exchange is limited to Exchange 2003 and 2007. SQL 2008 and 2008 R2 also support SCC.

2.7 Archive Vault Configuration Support

For each onQ Appliance that will host an onQ Archive Vault, an Archive Vault virtual machine (VM) will be present on that onQ Appliance. The Archive Vault VM, in turn, manages the Archive Vault Storage Modules (SM) attached to the onQ Appliance.

The following configurations are supported:

- Standalone Configuration
- DR-only Configuration

- Redundant Configuration
- Shared-Instance Configuration
- Unsupported Configurations

Related Topics

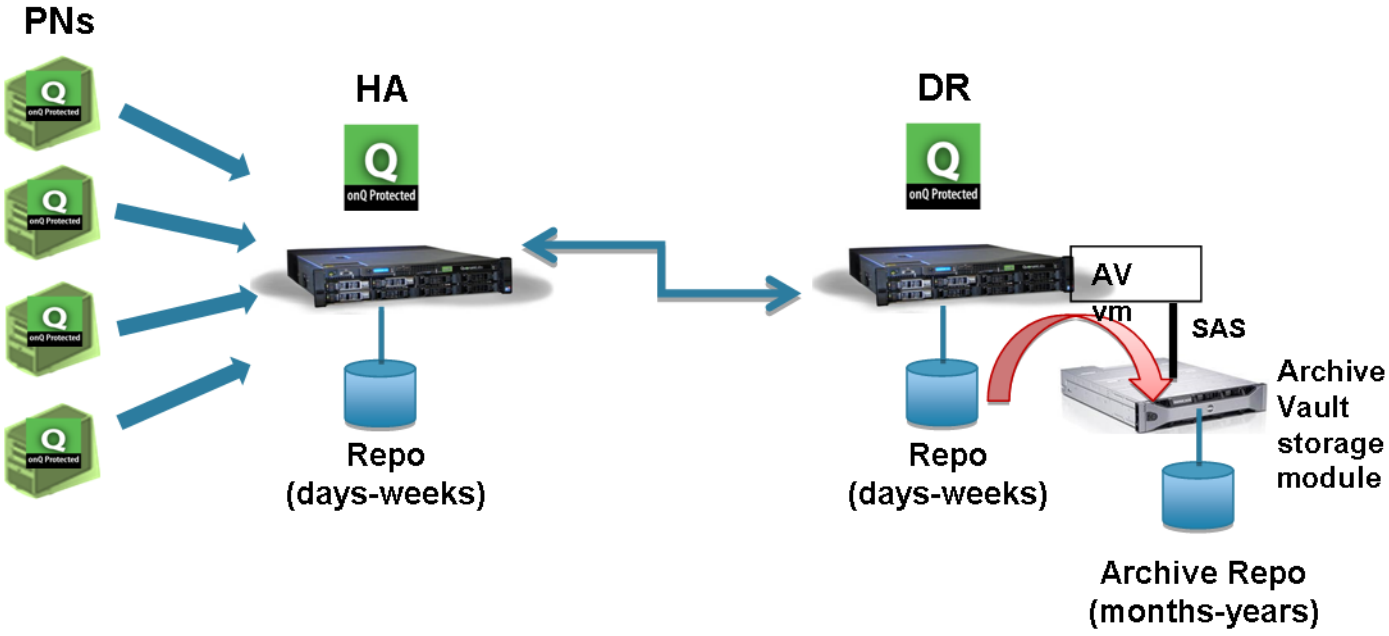
Unsupported Configurations

2.7.1 Standalone Configuration

In this configuration, the onQ Archive Vault isn't hosted by a DR or an HA - it is running on a standalone onQ Appliance (aka *Archive Vault Gateway*). Unlike all other configurations where the onQ Archive Vault is directly attached to a JBOD, this Archive Vault Gateway is connected to separate storage (for example, iSCSI, SAN) through a 10Gbps interface or a fiber channel interface card. The Archive Vault Gateway does require that an HA or DR be present elsewhere on the network.

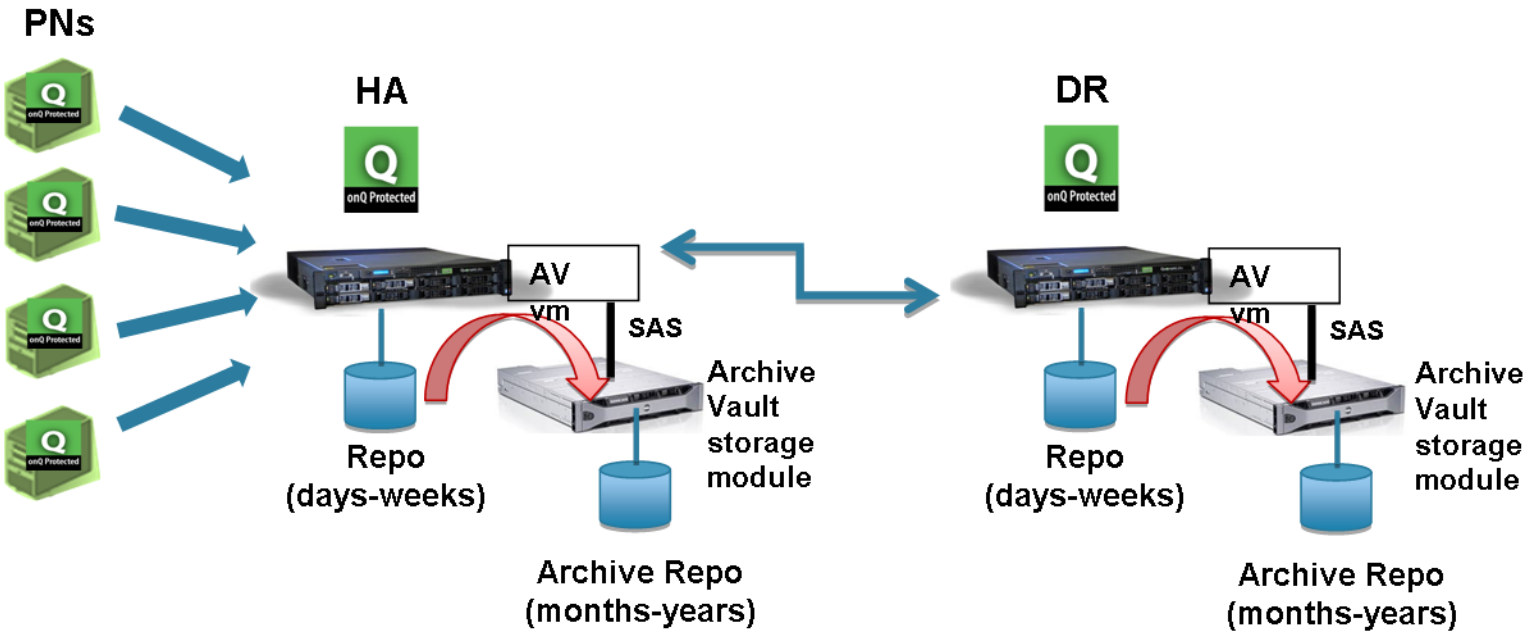
2.7.2 DR-only Configuration

In this configuration, the only instance of the onQ Archive Vault is attached to the onQ DR Appliance. Thus, there is a single instance of the long-term archive in a remote location. While protected by numerous fault-tolerant hardware features such as RAID6 disk drives and dual power supplies, this configuration might not satisfy some regulatory compliance standards, since there is no replication of long-term archive data.



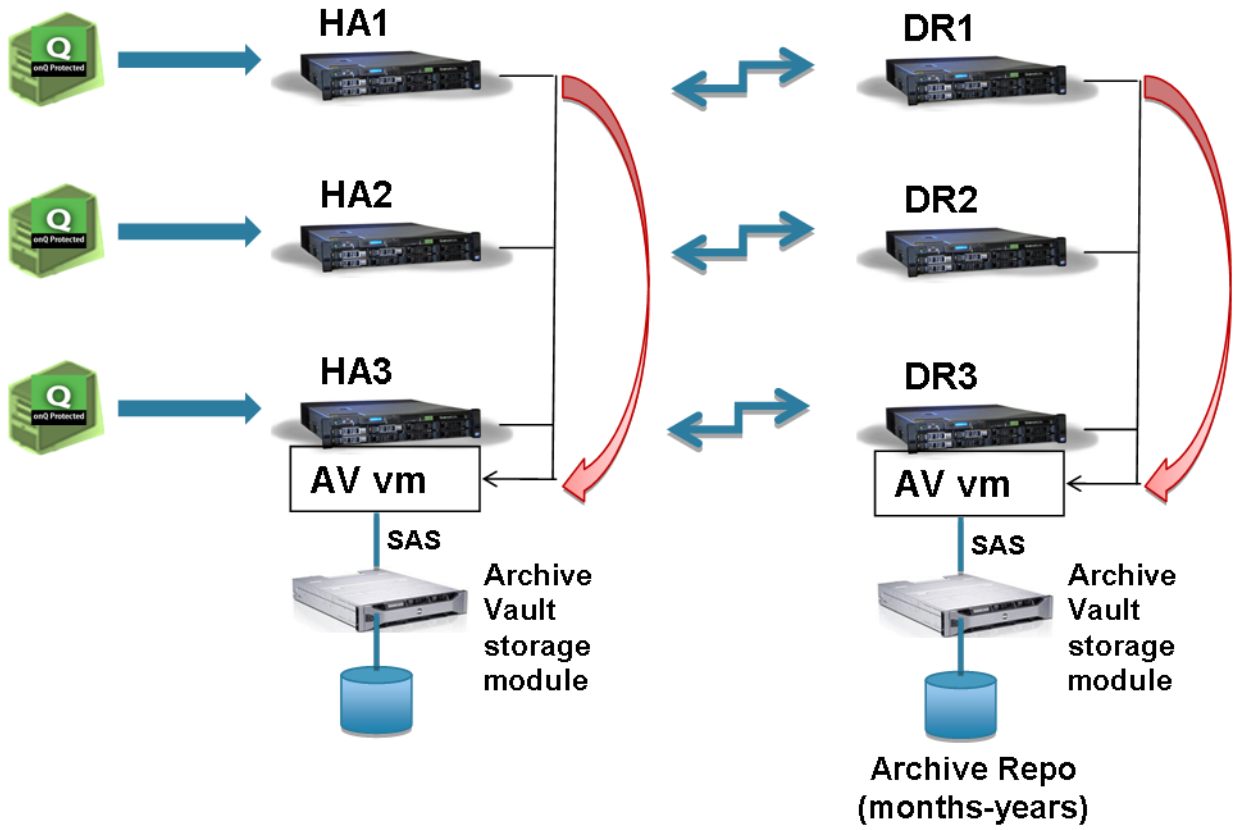
2.7.3 Redundant Configuration

In this configuration, there are instances of the onQ Archive Vault attached to both the HA Appliance and the DR Appliance. This configuration allows for both remote and redundant storage of long-term archive data, thus satisfying many regulatory compliance standards.



2.7.4 Shared-Instance Configuration

In this configuration, multiple HA Appliances or DR Appliances (or multiple tenants in a multi-tenant Appliance installation) are able to access a single onQ Archive Vault. This configuration allows for optimization of storage of long-term archive data on a single onQ Archive Vault even though multiple onQ Appliances might be present. The Shared-Instance configuration is also supported with DR Appliances exclusively, in situations where redundancy is not necessary.



2.7.5 Unsupported Configurations

An HA Appliance accessing an onQ Archive Vault that is hosted on a DR Appliance is not supported. Conversely, a DR Appliance accessing an onQ Archive Vault that is hosted on an HA Appliance is not supported.

3

Limitations

- User Interface Limitations
- Linux RN Limitations
- Windows Server Limitations
- Volume Mount Limitations
- CIFS Mount Limitations on Windows
- Bare Metal Restore Limitations
- File-level Restore Limitations
- Windows Cluster Services Limitations
- Agent-less PN Enrollment Limitations

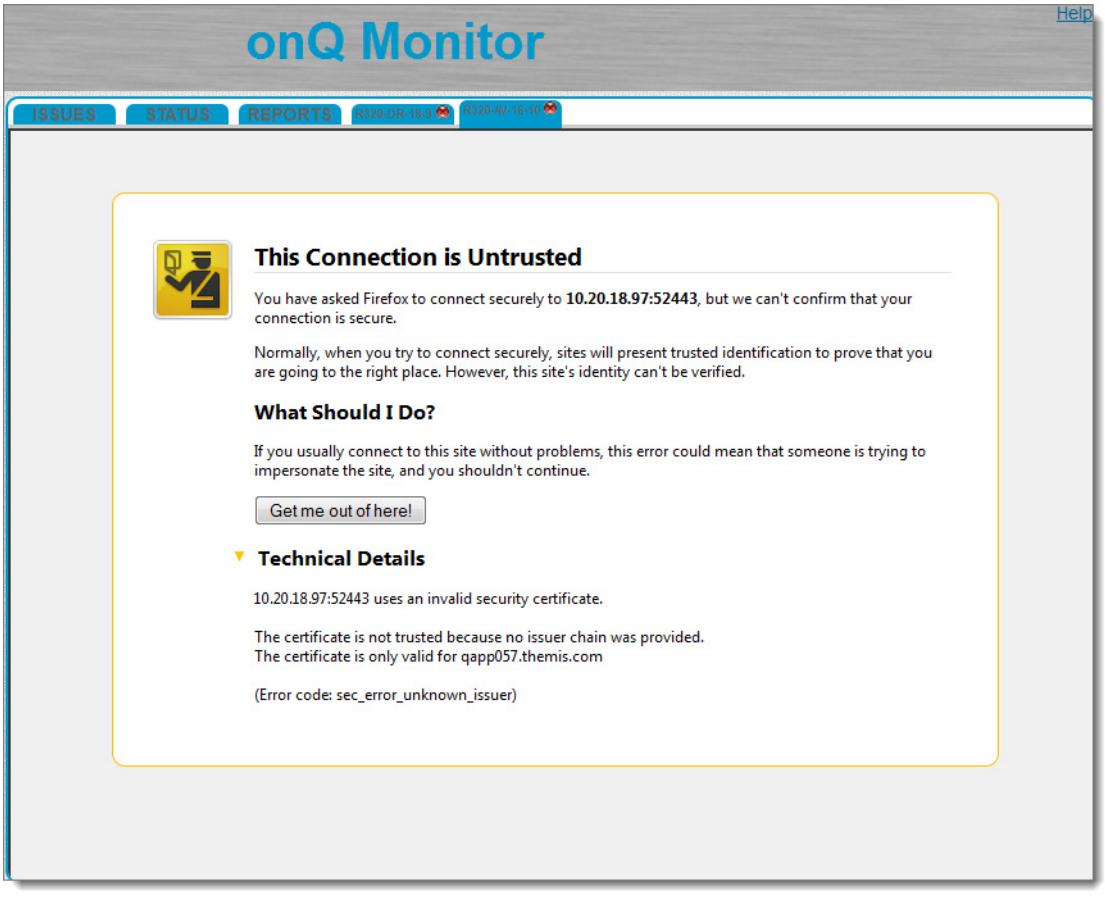
3.1 User Interface Limitations

- When you start an RN in test mode, the RN may not work with browsers running on clients that use non-US locale and keyboards. RNs in production mode and connected via RDP work as expected.
- Firefox 10.0-28.0 is not supported for enrollment (Protect Me button) on 64-bit operating systems because Firefox is unable to retrieve the node information that's required.
- When using older versions of some web browsers, you might need to install JRE (Java) 1.5 or above to use the RN console from the UI. For more information, go to Launch recovery node's console.
- You might need to modify security and popup settings for proper RN console operations from the UI.
- You must add the security certificate for the onQ Monitor to the permanent exception list using a separate browser window, not a browser tab; if you don't:

- the onQ Monitor's dashboard does not load and stops short of a certificate error, preventing you from launching the onQ Monitor's UI.
- onQ Monitor fails to return details of DR Appliance or AV Appliance: the hyperlink in the status page indicates `No Response from appliance` for DR Appliance and `Not Enrolled on AV` for AV Appliance.

Because the source URL for an iframe is not made visible, for security reasons, browsers do not allow you to accept unknown security certificates within those iframes.

To add the certificate to your exception list, Control-click the same link to open the page in a new browser *window* instead of a browser *tab*. The new page will be subject to pop-up blockers, but the page displays the normal range of browser options, enabling you to add the certificate to the permanent exception list. All future attempts to launch onQ Monitor's UI in a tab work as expected.



- When you launch the RN's console, it opens in a different browser window. If you browse to a different window, the console might lose connection after 60 seconds. In this case, you must launch the console again.

3.2 Linux RN Limitations

Recovery Nodes running Linux have the following limitations:

- To support booting of Recovery Nodes (RN) for RHEL 5.x/CentOS 5.x servers, you must download the appropriate version of Xen-compatible kernel packages and install them on the Protected Nodes (PN) before enrolling or attempting to build the RNs for such clients. The installed Xen kernel will not be an active kernel on the PN; instead, that kernel will be used to boot the RNs. For more information, go tothe Online Help "Install kernel-xen RPM package" in onQ Administrator's Guide.
- The filesystem format is required to be set in the Protection Node Configuration for each RN disk to ensure that onQ protection starts correctly.
- For RHEL 5.x/CentOS 5.x, the root disk partition (`/`) must be set to Format (file system type) parameter `ext3` for the RN to boot correctly. The enrollment procedure, outlined in the Online Help"(Agent-based Linux PNs) Enroll protected nodes" in onQ Administrator's Guide, accounts for this limitation.
- Linux PNs may return 256 code during an auto-upgrade. If you encounter this error, restart (stop/start) protection to re-initiate the auto-upgrade.
- Network Manager, if installed on any Linux PN, must be turned off for self-tests to work correctly. Similar action needs to be taken if Kudzu services are running on the RHEL 5.x/CentOS 5.x versions.
- BMR and incremental failback for Linux nodes are not supported in this release.

3.3 Windows Server Limitations

Consider the following limitations if any RN is running Windows Server:

- (*Windows Server 2012 R2 Only*) It is possible that the mouse may not be activated when an RN is running in test mode via Java console. Attempt to run the RN in production mode, and use an RDP session to access the RN. This bug will be resolved pending a XenServer release of a new BIOS or Microsoft fix.
- (*Windows Server 2012 & R2*) When a REFS volume in Windows Server 2012 or R2 is protected, the Recovery Node for the same will be in standard NTFS. This is also true of a Bare Metal Restore.
- (*Windows Server 2012 & R2*) When a deduplicated volume in Windows Server 2012 R2 is protected, the Recovery Node for the same will again be a standard NTFS. This is also true of a Bare Metal Restore. Additional caution should be used in allocating RN disk space as the Quorum agent on the Protected Node will report actual disk usage, which might be far less than when data is stored in its normal form.
- (*All Windows Server platforms*) When onQ event log reports 3010 code for any Windows PNs, both 32-bit and 64-bit, the auto-upgrade fails. Perform *one* of the following workarounds in the order that they are listed:
  - (Agent-based PNs) Restart onQ Service on the PN.
  - Reboot the PN.
  - Uninstall and reinstall the latest msi package at `C:\Program Files\Quorum\installer`.

3.4 Volume Mount Limitations

- Volumes only reachable through a GUID path are not supported.
- Backup source arguments that aren't drive letters or volume mount paths that begin with a drive letter are not supported.
- Paths as backup source arguments to onQ that are not volumes or volume mount points are not supported when using with VSS.
- Volume mount points cannot contain the characters outside of A-Z, a-Z, 0-9, underscore and a dash. Volume mount points with spaces are not supported.
- Total string length of drive letter mounts and volume mount points, including separators in excess of 250ch/s, are not supported.
- Volume mount point that are based on a volume that has other mount points that refer to it (multiple mount points) are not supported.
- Volume mount points used to protect MS cluster environment on onQ are not supported.
- Volume mount points might not be visible on the RN under certain circumstances if the folder that contains the mount points is hidden or in a system folder. You might need to set the folder attributes to display system/hidden folders so as to display mount points.
- A maximum of 15 logical drives can be protected per PN.

3.5 CIFS Mount Limitations on Windows

- In order for onQ to back up shared volumes, there are some prerequisite steps that you need follow as outlined in the Online Help"Enable onQ to back up shared volumes" in onQ Administrator's Guide.
- When file objects are restored from a NAS backup, the kernel security objects may not be written. The data contents of the file will be restored as it was backed up.
- You cannot specify the size of a vdisk when you create or update the vdisk. This limitation is true for both automatic enrollment and manual enrollment.

- You cannot perform a BMR for CIFS drives. BMR continues to support only fixed drives on protected Windows systems.
- Network shared volumes can be extremely large; therefore, restoring to a specific PN is not an efficient way to retrieve data. Instead, restore the data to a NAS appliance. A BMR can only restore to local drives.
- Recovery Nodes will not include data for a network shared drive. However, Recovery Nodes started in production mode can connect to the configured remote drives.
- When an RN boots as part of a self-test or because you choose to boot it in test mode, the remote mounts will not be accessible.
- Backup performance depends on your PN's and NAS's bandwidth, I/O limitations, and other resources.
- File Backup Only option cannot be used with a remote drive.
- See the CIFS limitation outlined in Agent-less PN Enrollment Limitations.

3.6 Bare Metal Restore Limitations

- When you perform a failback to restore a VMware-based protected node to a dissimilar hardware or virtual environment, Windows might require reactivation.
- BMRs that are performed to the same or new hardware require that you upload specific drivers for any disk, network adapter, and any other hardware.
- After an incremental failback to a non-Xen PN (physical PN, ESX PN, Virtual PC PN), users are required to reinstall drivers.
- For a Windows 2012 server, a REF's partition/volume will be restored as standard NTFS.
- For a Windows 2012 server, if you choose to restore a volume with deduplication enabled, the data will be restored to its full size. Simply enable deduplication after the BMR completes.
- See the dynamic disk limitation outlined in Agent-less PN Enrollment Limitations.
- When adding each volume in the disk configuration interface, you might need to wait a short period of time while the volumes are being created.

- There is a maximum allowed duration of 72 hours for a restore to complete. Exceeding 72 hours might result in a reboot of QUARK and a loss of the restored image to that point.

Related Topics

File-level Restore Limitations

3.7 File-level Restore Limitations

- (Agent-based PNs) Restoring using the `Original location of files` **Target** option is not supported for:
  - Linux.
  - From DR or AV.
  - From a volume on source PN that doesn't exist on the target server. For example: Files from `G:\` volume on `PN1` cannot be restored to `PN2` where `G:\` doesn't exist.
- (Agent-less PNs) FLR is supported; however, you must use the Other option because an FLR to same or different PN relies on an agent running on the target PN. You can combine the `Other` **Destination PN to restore** option with the `Original location of files` **Target** option.

3.8 Windows Cluster Services Limitations

- Suppose cluster nodes A and B are configured to use the same disks for their operation. If both of them are enrolled for protection, along with their shared disk configuration, only the cluster node that is currently using the shared disks can be backed up. If you attempt to back up the cluster node that does not have access to the shared disks, the backup will fail. Because the IP address of all the cluster nodes in the host table are set to the IP address of the cluster, any passive node in the cluster will indicate a RED status and backups will fail.
- When you start a cluster RN, you must select the RN for the node that was active at the time your desired backup was generated. For example, if you want the most recent data, then select the node with the most recent backup.The RN for the other nodes will contain stale information. If you

want to roll back to a previous backup, you must select the RN that was active at that point in time.

- Recovery Node disks used in this context are basic disks, each with a maximum size limit of 2TB. The logical volume size for these disks is not expandable.
- If the PN configuration is switched between a standard and a cluster, the RN will be reinitialized and rebuilt.
- A cluster node that is protected with this configuration whose RN is brought in production can only be configured to work with virtual SAN. Registry modifications are done in the RN from the information provided in the onQ Portal for the specific node.
- A cluster node protected with this configuration is not suitable for a Bare Metal Restore as cluster service configurations are modified on the RN to work with a virtual SAN. Also, the failover cluster node is evicted from the configuration while preparing the RN. Only application data (SQL database, for example) can be retrieved/restored.
- There are additional cluster limitations if the PNs are agent-less. Go to Agent-less PN Enrollment Limitations.

3.9 Agent-less PN Enrollment Limitations

- Whether you use agent-less or agent-based enrollment, onQ can perform three PN backups simultaneously. However, with agent-less enrollment, your agent-less PNs must reside on different ESX/ESXi hosts.
- Clustering, as outlined below, is supported for agent-less PNs; however, a shared network and shared IP is not supported.
- Due to a VMware limitation, dynamic disks are not supported. Do not attempt to protect them. If you protect them, the volumes and mount points based on the dynamic disks will not be available in the RN. Also, you will not be able to use QUARK to restore the dynamic disks.
- For agent-less PNs, FLR is supported; however, there are limitations. Go to File-level Restore Limitations.

- VMware does not support snapshots of raw disks, RDM physical mode disks, or guest operating systems that use an iSCSI initiator in the guest. As outlined above, snapshots for dynamic disks and logical drives are not supported. Therefore, Quorum protection for agent-less PNs supports cluster nodes and local drives so long as they are primary partitions on basic disks. If the RN of such a cluster node is brought up in production mode, then the RN will connect to a production SAN as configurations from the PN will be replicated and are available on the RN.
- CIFS volumes are not supported. When a NAS share is mounted on a PN, the security credentials are accessible for applications running in the PN. However, in the case of a PN proxy, those security credentials aren't accessible.
- A virtual machine hosted on ESXi servers that are managed by vMotion can be enrolled as agent-less PNs and protected by onQ; however, if you then migrate that virtual machine to a different host using vMotion, onQ is unable to back up that virtual machine.
- Using the onQ Filter Driver, as described in (Agent-based PNs) Methods of Performing Incremental Backups, to back up agent-less PNs is not supported.
- If you start an agent-less PN's RN in production mode, the DR Appliance will not be able to communicate with the HA Appliance even after you install the onQ Service on the RN. In the event that you need to back up the RN because the PN failed, start the RN on HA Appliance, install onQ Service on the RN, then continue to back up the RN. Alternatively, change DR Appliance's role to HA, then run the RN in production mode.
- If a PN on the HA has custom (RN) networks defined and the DR has the default network defined, the RN on that DR might be configured with the custom networking information instead of the default networking information. The workaround is to create an `xvf.dat` file, then save it to `C:\Program Files\Quorum\usr` on the HA's PN. As in the example below, edit the file to include the PN's networking information. In the following example, `10.20.16.144` is the default IP of the PN enrolled on the HA, `255.255.248.0` is the default subnet mask, and `10.20.16.1` is the default gateway. Afterward, rebuild the RN. The PN on the HA

continues to show the custom networks while the RN on the DR indicates what you specified in the `xvf.dat` file.

```
@@XV_PN_IP
10.20.16.144
@@XV_PN_MASK
255.255.248.0
@@XV_PN_GW
10.20.16.1
```

Related Topics

“(Agent-less PNs) Connection and Backup Problems” in onQ Administrator’s Guide

4

Updates

- Upgrade Support
- Update Appliance software

4.1 Upgrade Support

Upgrading *directly* to version 3.8 SP1 is supported for onQ Appliances and protected nodes running 3.8 HF1.

When upgrading, take note of the following:

- When you upgrade your onQ Appliance or AV to 3.8 SP1, if onQ's or AV's timezone is set to anything other than United States time, the installer resets the timezone such that it shows the incorrect timezone for Backup Status/Archival Status and onQ Status/AV Status. (The timezone in the dialog's drop-down list is correct.) To set the Backup Status/Archival Status and onQ Status/AV Status to the correct timezone:
  1. Stop protection.
  2. Open the **Modify Local onQ Setup**/**Modify Archive Vault Setup** and save the expected timezone: (**APPLIANCE CONFIG** | **ARCHIVE CONFIG** tab > **onQ (LOCAL)** | **ARCHIVE VAULT** page > **MODIFY** button > **SAVE** button).
  3. Reboot the onQ and AV.
- After an onQ upgrade to 3.8 SP1, you must reinstall the Linux PN proxy and the Windows PN proxy on each ESX/ESXi host enrolled with the onQ. This reinstall enables improvements to the PN proxy's handling of power failures. Perform the following steps for each ESX/ESXi host enrolled with the onQ:
  1. Log on to the HA Appliance's onQ Portal.
  2. Click the **PROTECTION CONFIG** tab.

3. Click the double-plus button (**++**).

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

4. Provide the root userid and password to access the ESXi host.
5. Select **Re-install Windows proxy** or **Re-install Linux proxy**, depending on the proxy you're trying to reinstall, then **ENROLL**.
6. Observe the messages that appear in the pop-up window to verify that the proxy reinstalled successfully.

- After an onQ upgrade, you must refresh your browser using CTRL-F5. If you don't, portal changes will not be activated and some user actions will not work.
- An onQ upgrade to 3.8 SP1 does not involve any major repository migration or restructuring.
- You must manually reboot the onQ Appliance after the onQ upgrade completes.
- Upgrade packages are deployed through onQ Central. onQ and AV will auto-download the necessary onQ packages when the daily cron job runs. This job runs between 10pm and 2am each day. This deployment occurs for all onQ Appliances for which onQ Central is set to allow download of the latest software packages.

4.2 Update Appliance software

Quorum regularly releases new versions of onQ Appliance software, including service packs and, from time to time, hot fixes.

Upgrade packages are deployed through onQ Central. onQ and AV will auto-download the necessary onQ packages when the daily cron job runs. This job runs between 10pm and 2am each day. This deployment occurs for all onQ Appliances for which onQ Central is set, by Quorum Support, to allow download of the latest software packages.

To *automatically* apply an update:

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

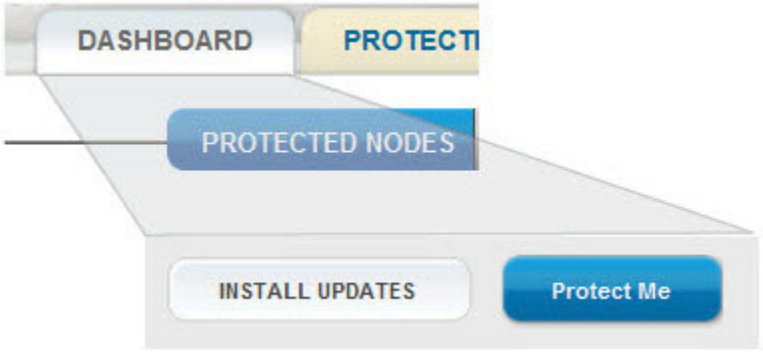
- Unless instructed to do so in , there is no need to stop protection to apply the updates. If protection is on, onQ quiesces the system, stops protection, applies the new software, then restarts protection. The upgrade process reboots the onQ, if required.
- For onQ Appliances that are multi-tenant, you must upgrade them one tenant at a time and in one session. Unless instructed to do so in *(Start Here) Upgrade Support and Requirements*, it is not necessary to reboot the onQ Appliance after each tenant upgrade. A reboot after the last upgrade will suffice.
- You must upgrade both HA Appliances, DR Appliances, and onQ Archive Vaults synchronously.

Before you begin, verify that your upgrade path is supported and learn about any pre-upgrade and post-upgrade requirements.

1. Log on to the onQ Portal. You must log on as `varadmin` user.
2. Trigger the download of the onQ package, if it hasn't already downloaded.
   a. Go to **APPLIANCE CONFIG** tab > **ADVANCED** button > **Software Upgrades** page.
   b. Click the **Check for Updates** button.

   Mouse-over the Check for Updates button to display the status of the download.

   The onQ Portal's **DASHBOARD** displays an **Install Updates** button after the new software downloads from onQ Central. The **onQ STATUS** page > **Issues** pane also displays an `Appliance software upgrade available` message.



3. Apply the new software:

a. Go to the **DASHBOARD** tab.

b. Scroll down and click the **INSTALL UPDATES** button.

c. Click the **UPGRADE** button when prompted to initiate the upgrade.

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

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

   Immediately after this onQ Appliance upgrade, onQ automatically initiates a PN software (agent-based) and PN proxy software (agent-less) upgrade as outlined in:

   - *(Agent-based Windows PNs) Update node software*
   - *(Agent-based Linux PNs) Update node software*
   - *(Agent-less Windows/Linux PNs) Update PN Proxy software*