

# Features - Quick Recovery Agent

---

## TABLE OF CONTENTS

### OVERVIEW

### SYSTEM REQUIREMENTS - QUICK RECOVERY AGENT

### INSTALLATION

- Install the Quick Recovery Agent - Windows
- Install the Quick Recovery Agent - Unix

### RECOVERY POINTS FOR THE QUICK RECOVERY AGENT

### CREATING A QR VOLUME

### RECOVERING QR VOLUMES

### RESTORE DATA - SNAPVAULT

### RESTORE DATA - ONTAP SNAPSHOT

### CONFIGURATION

- Subclients - Quick Recovery Agent
- QR Policies
- Scratch Volume Pools

### MANAGEMENT

- QR Volume Creation History
- QR Volume Recovery History

### USE CASES

- QR Disaster Recovery Solution for Building a UNIX Standby Oracle Server with Log Recovery
  - QR Disaster Recovery Solution for Building a UNIX Standby Oracle Server
  - QR Disaster Recovery Solution for Building a Windows Standby Oracle Server
  - QR Disaster Recovery Solution for Building a Windows Standby SQL Server
  - QR Disaster Recovery Solution for Building a Windows Standby Exchange Server in Quick Implementation Mode
  - QR Disaster Recovery Solution for Building a Windows Standby Exchange Server in Quick Recovery Mode
  - QR Disaster Recovery Solution for an Exchange Cluster Using a Standby Server
  - QR Disaster Recovery Solution for Building a Standby SQL Server on a MS Cluster
-

# Overview - Quick Recovery<sup>®</sup> Agent

---

This feature/product/platform is on Extended Support in this release. See [Deprecated Features, Products, and Platforms](#) for more information.

Choose from the following topics:

- Introduction
- Supported Data Types
- Tree Levels in the QR<sup>™</sup> Agent
- License Requirement
- Snapshot Support
  - Windows Server 2003 Considerations
- LAN Copy Manager
- Application iDataAgent Integration
- Disaster Recovery Considerations
- User Cases - Disaster Recovery Solutions for Building Standby Servers:
  - UNIX Standby Oracle Server with Log Recovery
  - UNIX Standby Oracle Server
  - Windows Standby Oracle Server
  - Windows Standby SQL Server
  - Windows Standby Exchange Server in Quick Implementation Mode
  - Windows Standby Exchange Server in Quick Recovery Mode

## See Also

- QR<sup>™</sup> Policies
  - Scratch Volume Pools
  - System Storage Model
- 

## INTRODUCTION

The Quick Recovery<sup>®</sup> Agent works in conjunction with a snapshot engine to create QR Volumes. A QR Volume is a block-level copy of the primary disk volume that contains your application data. Each QR Volume captures the state of a primary volume at a particular point in time — the time when the snapshot was created.

The Quick Recovery Agent builds on the existing infrastructure of the software to provide control and management of hardware and software-based volume snapshot technology, along with a framework for disk-to-disk block-level volume copies via SAN or LAN-based data movement. The Quick Recovery Agent integrates with major storage-intensive applications such as Microsoft SQL Server and Microsoft Exchange 2000 to ensure that the application data objects are properly synchronized and easily recovered. Through this combination of snapshot technology, volume replication, and application awareness, the Quick Recovery Agent provides an unprecedented form of data protection with rapid recovery from storage-related failures. When installed in conjunction with iDataAgents for backup and recovery, the Quick Recovery Agent forms part of a total storage management solution.

Multiple point-in-time QR Volume copies of a given primary volume may be created, provided that sufficient scratch pool volumes are available to receive the data. Also, optionally, a single QR Volume may be created with the incremental update feature. When incremental update is enabled, the Quick Recovery Agent will periodically perform an incremental, block-level update of the QR Volume with the data that has been changed on the primary disk since the previous update. This process is much more efficient than a full volume copy.

---

## USE OF RECOVERY POINTS

Combined with the QSnap<sup>®</sup> snapshot enabler, the Quick Recovery<sup>®</sup> Agent can be used to create and maintain Consistent Recovery Points that preserve the states of QR Volumes after incremental updates. These point-in-time copies of your volumes can be mounted, shared on a network, or copied back, providing another layer of protection and flexibility for your backup solution. See [Recovery Points for the Quick Recovery Agent](#) for more information about how the QSnap enabler and Quick Recovery Agent can work together.

---

## APPLICATION AWARENESS

Application data that has been lost or modified can be recovered simply by mounting the QR Volume and accessing the data directly. When configured to protect and recover an application volume, the QR Agent will intelligently quiesce the application (if necessary). For a complete listing of applications supported



To perform a data protection operation using this Agent a specific Product License must be available in the CommServe® Server.

Review general license requirements included in License Administration. Also, View All Licenses provides step-by-step instructions on how to view the license information.

[Back to Top](#)

---

### SNAPSHOT SUPPORT

The Quick Recovery® Agent is designed to work in conjunction with the following snapshot engines, which provide snapshot functionality for data protection operations:

- File System Snapshot
- QSnap® Enabler for the Quick Recovery Agent
- VSS for the Quick Recovery Agent
- EMC SnapView
- Hitachi HDS QuickShadow
- Hitachi HDS ShadowImage
- Hitachi HDS TrueCopy
- ONTAP SnapVault for the Quick Recovery Agent
- OSSV SnapVault for the Quick Recovery Agent

In addition, the use of Generic Enablers allows the Quick Recovery Agent to employ third-party (hardware) snapshot capabilities.

For more information on the snapshot feature, refer to Snapshots. For more detailed snapshot support information, refer to Support Information - Snapshot Engines.

---

### WINDOWS SERVER 2003 CONSIDERATIONS

Consider the following factors when choosing between VSS and QSnap for your snapshot engine on Windows Server 2003:

- VSS is the default enabler for Windows Server 2003.
- QSnap enabler works on Windows Server 2003 as well as on other platforms.
- VSS does not quiesce the server when taking a snapshot; QSnap does quiesce the server.

[Back to Top](#)

---

### LAN COPY MANAGER

Although the Quick Recovery® Agent is optimized for use with a SAN, it may still be used effectively in a more traditional LAN environment. In this case, the LAN copy manager facility can be used to transfer data between source and destination volumes. The LAN copy manager is especially useful for copying volumes to remote sites and thus can function as a WAN copy manager as well. For more information, see Copy Managers.

[Back to Top](#)

---

### APPLICATION /DATAAGENT INTEGRATION

The QR™ Agent can work together with the following iDataAgents to provide an integrated data protection and recovery system for your database volumes:

- SQL Server iDataAgent
- Exchange Database iDataAgent
- Oracle iDataAgent (Unix and Windows)

In this scenario, the QR Agent creates and maintains a copy, or QR Volume, of the database and log volumes. The application iDataAgent is then used to back up the associated application's transaction logs.

During the recovery process, the QR Agent will recover the application volumes to their state as of the last update of the QR Volume. Transaction logs that were backed up after that point in time are then restored and applied (replayed) by the application iDataAgent.

This configuration takes advantage of both the speed of recovery offered by the QR Agent, and the efficiency of transaction log backups (which are typically less resource intensive than database backups and can be scheduled to occur frequently).

- The log recovery using the QR Agent is not supported with Exchange.
- It is recommended that you do not change the name of the client during installation as this can cause the

application discovery to fail.

- During **Add App**, SQL Backward compatibility Tool should be installed in order to detect SQL databases created using higher version for SQL.

[Back to Top](#)

---

## **DISASTER RECOVERY CONSIDERATIONS**

- Before you use your agent, be sure to review and understand the associated full system restore (or disaster recovery) procedure. The procedure for some agents may require that you plan specific actions or consider certain items before an emergency occurs. See Disaster Recovery for more information regarding your agent.

[Back to Top](#)

---

## System Requirements - QuickRecovery Agent

This feature/product/platform is deprecated in this release. See [Deprecated Features](#), [Products](#), and [Platforms](#) for more information.

The following requirements are for the QuickRecovery Agent:

If using this Agent with QSnap, refer to the following for additional support information: [System Requirements-QuickRecovery Agent - Snapshots Support](#) for information on the operating system vendors supported for QSnap with this Agent.

OPERATING SYSTEM		PROCESSOR	
<b>AIX</b>	AIX 5.3 64-bit with technology level 6 and runtime library xIC.rte 8.0.0.0 or higher	Power PC (Includes IBM System p)	
<b>HP-UX</b>	HP-UX 11i v2 (11.23) 64-bit	PA-RISC	
	HP-UX 11i v2 (11.23) 64-bit (Supported with the Generic Enabler only.)	Itanium	
	HP-UX 11i v1 (11.11) 64-bit with OS patch PHCO29328 (contact Hewlett Packard to obtain the patch)	PA-RISC	
	HP-UX 11i v1 (11.11) 32-bit with OS patch PHCO29328 (contact Hewlett Packard to obtain the patch)	PA-RISC	
<b>LINUX</b>	<b>RED HAT ENTERPRISE LINUX/CENTOS</b>		
	Red Hat Enterprise Linux/CentOS 5.x with glibc 2.5.x	Intel Pentium or compatible minimum required	
	Red Hat Enterprise Linux/CentOS 5.x with glibc 2.5.x	x64	
	Red Hat Enterprise Linux/CentOS 4.x with a minimum of glibc 2.3.4	x64	
	Red Hat Enterprise Linux/CentOS 4.x with a minimum of glibc 2.3.4	Intel Pentium or compatible minimum required	
	<b>SUSE LINUX (SLES)</b>		
	SuSE Linux 10.x with glibc 2.4.x	Intel Pentium or compatible minimum required	
	SuSE Linux 10.x with glibc 2.4.x	x64	
	<b>SOLARIS</b>	Solaris 9 64-bit with Service Packs 111711-02 and 111712-02	Sparc5 or higher recommended
		Solaris 10.x with a minimum of SunOS (Sparc) Patch 119963-14	Sparc5 or higher recommended
<b>WINDOWS</b>	<b>WINDOWS 2008</b>		
	Microsoft Windows Server 2008 32-bit and x64 Editions* *Core and R2 Editions not supported	All Windows-compatible processors supported	
<b>WINDOWS</b>	<b>WINDOWS 2003</b>		
	Microsoft Windows Server 2003 32-bit and x64 Editions with a minimum of Service Pack 1	All Windows-compatible processors supported	

### CLUSTER - SUPPORT

The software can be installed on a Cluster if clustering is supported by the above-mentioned operating systems.

For information on supported cluster types, see [Clustering - Support](#).

### HARD DRIVE

#### AIX, HP-UX, LINUX, AND SOLARIS

225 MB minimum of hard drive space for software

100 MB of additional hard disk space for log file growth

10 MB of temp space required for install or upgrade (where the temp directory resides)

---

## WINDOWS

101 MB minimum of hard disk space for software/ 148 MB recommended  
 50 MB of additional hard disk space for log file growth  
 725 MB of temp space required for install or upgrade (where the temp folder resides)

## MEMORY

---

## AIX, HP-UX, LINUX AND SOLARIS

---

### WINDOWS

16 MB RAM minimum required beyond the requirements of the operating system and running applications  
 Swap space = 2\*RAM size  
 32 MB RAM minimum required beyond the requirements of the operating system and running applications

### PERIPHERALS

DVD-ROM drive  
 Network Interface Card

### MISCELLANEOUS

---

### NETWORK

TCP/IP Services configured on the computer.  
 The operating system must have been installed with at least the `user level software` option selected.

---

### MICROSOFT VISUAL C++

Microsoft Visual C++ 2008 Redistributable Package is automatically installed. Note that Visual C++ 2008 Redistributable Package can co-exist with other versions of this software.

---

### NET FRAMEWORK

.NET Framework 2.0 is automatically installed. Note that .NET Framework 2.0 can co-exist with other versions of this software.

---

### SELINUX

If you have SELinux enabled on the client computer, create the SELinux policy module as a root user before performing a backup. The SELinux Development package must be installed on the client.

To create an SELinux policy module, perform the following steps as user "root":

1. Create the following files in the `/usr/share/selinux/devel` directory:

File Name	Content of the File
<code>&lt;directory&gt;/&lt;file_name&gt;.te</code> where: <code>&lt;directory&gt;</code> is <code>/usr/share/selinux/devel</code> <code>&lt;file_name&gt;</code> is the name of the Unix file, created to save the policy module statement. It is a good idea to use the same name for policy module and the file. For example: When you are creating a policy module for backup_IDA application, you can use the following file name: <code>backup_IDA.te</code>	The content of the file should be as follows: <code>policy_module(&lt;name&gt;,&lt;version&gt;)</code> <code>#####</code> where: <code>&lt;name&gt;</code> is the name of the policy module. You can give any unique name to the policy module, such as a process or application name. <code>&lt;version&gt;</code> is the version of the policy module. It can be any number, such as 1.0.0. For Example: While creating a policy module for the backup_IDA application, you can use the following content. <code>policy_module(backup_IDA,1.0.0)</code>
<code>&lt;directory&gt;/&lt;file_name&gt;.fc</code> where: <code>&lt;directory&gt;</code> is <code>/usr/share/selinux/devel</code> <code>&lt;file_name&gt;</code> is the name of the Unix file, created to save the policy module statement. It is a good idea to use the same	The content of the file should be as follows: Note that the following list of files is not exhaustive. If the process fails to launch, check <code>/var/log/messages</code> . Also, if required, add it to the following list of files. <code>/opt/&lt;software installation directory&gt;/Base/libCTreeWrapper.so -- gen_context(system_u:object_r:texrel_shlib_t,s0)</code>

name for policy module and the file.

For example: When you are creating a policy module for backup\_IDA application, you can use the following file name: backup\_IDA.fc

```
/opt/<software installation directory>/Base/libCVMAGuiImplgso -- gen_context
(system_u:object_r:texrel_shlib_t,s0)
/opt/<software installation directory>/Base/libdb2locale.so.1 -- gen_context
(system_u:object_r:texrel_shlib_t,s0)
/opt/<software installation directory>/Base/libdb2osse.so.1 -- gen_context
(system_u:object_r:texrel_shlib_t,s0)
/opt/<software installation directory>/Base/libDb2Sbt.so -- gen_context
(system_u:object_r:texrel_shlib_t,s0)
/opt/<software installation directory>/Base/libdb2trcapi.so.1 -- gen_context
(system_u:object_r:texrel_shlib_t,s0)
/opt/<software installation directory>/Base/libDrDatabase.so -- gen_context
(system_u:object_r:texrel_shlib_t,s0)
/opt/<software installation directory>/Base/libIndexing.so -- gen_context
(system_u:object_r:texrel_shlib_t,s0)
/opt/<software installation directory>/Base/libSnooper.so -- gen_context
(system_u:object_r:texrel_shlib_t,s0)
```

2. Create the policy file from command line. Use the following command. Ensure that you give the following commands in the /usr/share/selinux/devel directory.

```
[root]# make backup_IDA.pp
Compiling targeted backup_IDA module
/usr/bin/checkmodule: loading policy configuration from tmp/backup_IDA.tmp
/usr/bin/checkmodule: policy configuration loaded
/usr/bin/checkmodule: writing binary representation (version 6) to tmp/backup_IDA.mod
Creating targeted backup_IDA.pp policy package
rm tmp/backup_IDA.mod tmp/backup_IDA.mod.fc
[root]# semodule -i backup_IDA.pp
[root]#
```

3. Execute the policy module. Use the following command:

```
[root]# restorecon -R /opt/<software installation directory>
```

SELinux is now configured to work with this application.

#### DISCLAIMER

Minor revisions and/or service packs that are released by application and operating system vendors are supported by our software but may not be individually listed in our System Requirements. We will provide information on any known caveat for the revisions and/or service packs. In some cases, these revisions and/or service packs affect the working of our software. Changes to the behavior of our software resulting from an application or operating system revision/service pack may be beyond our control. The older releases of our software may not support the platforms supported in the current release. However, we will make every effort to correct the behavior in the current or future releases when necessary. Please contact your Software Provider for any problem with a specific application or operating system.

Additional considerations regarding minimum requirements and End of Life policies from application and operating system vendors are also applicable



# Install the Quick Recovery Agent - Windows

## TABLE OF CONTENTS

### Install Requirements

#### Before You Begin

#### Install Procedure

- Getting Started
- Cluster Selection
- Select Components for Installation
- Configuration of Other Installation Options
- Download and Install Latest Packs
- Client Group Selection
- Schedule Automatic Update
- Configure QSnap
- Verify Summary of Install Options
- Install Remaining Cluster Nodes
- Setup Complete

#### Post-Install Considerations

## INSTALL REQUIREMENTS

The following procedure describes the steps involved in installing the Quick Recovery Agent and QSnap on both cluster and non-cluster environment.

The Quick Recovery Agent is installed on the server for which you will be creating Quick Recovery volumes. (This computer is referred to as the *Client* computer in this install procedure.) In addition, you should also install any snapshot enablers you plan to use with the Agent.

Verify that the computer in which you wish to install the software satisfies the minimum requirements specified in System Requirements - Quick Recovery Agent.

Review the following Install Requirements before installing the software:

---

### GENERAL

- Review Install Considerations before installing the software.
- Agents should be installed only after the CommServe and at least one MediaAgent have been installed in the CommCell. Also, keep in mind that the CommServe® software and MediaAgent must be installed and running (but not necessarily on the same computer), before you can install the Agent.
- Close all applications and disable any programs that run automatically, including anti-virus, screen savers and operating system utilities. Some of the programs, including many anti-virus programs, may be running as a service. Stop and disable such services before you begin. You can re-enable them after the installation.
- Ensure there is an available license on the CommServe software for the Agent.
- Verify that you have the Software Installation Disc that is appropriate to the destination computer's operating system.

---

### CLUSTER SPECIFIC

- If you will be using QSnap, you must install QSnap on all the physical nodes in the cluster.
- Once the appropriate software has been installed onto the physical nodes, the Quick Recovery Agent and any snapshot enablers (except QSnap) can be installed from the active node in the cluster group using the following procedure. The software can also be automatically installed on all available passive nodes when it is installed in the cluster group, or you can choose to install any passive node(s) separately.
- Check the following on the cluster computer in which you wish to install the software:
  - Cluster software is installed and running.
  - Active and passive nodes are available.
  - Disk array devices configured with access to the shared array.
  - Public Network Interface Card is bound first, before the private Network Interface Card. (Does not apply to NetWare Cluster.)

## BEFORE YOU BEGIN

- Log on to the client as the local Administrator or as a member of the Administrators group on that computer.
- On a clustered computer, ensure that you are logged on to the **active node** as the Domain User with administrative privileges to all nodes on the cluster.

## INSTALL PROCEDURE

---

### GETTING STARTED


1. Place the Software Installation Disc for the Windows platform into the disc drive.  
After a few seconds, the installation program is launched.

If the installation program does not launch automatically:

- Click the **Start** button on the Windows task bar, and then click **Run**.
- Browse to the installation disc drive, select **Setup.exe**, click **Open**, then click **OK**.

**NOTES**

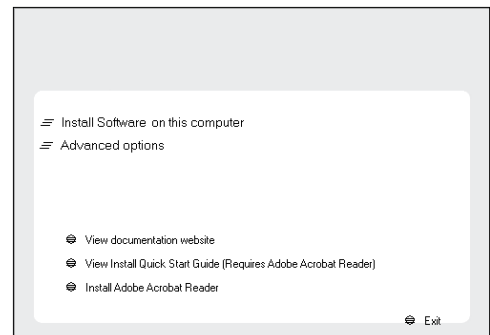
- If you are installing on Windows Server Core editions, mount to Software Installation Disc through command line, go to the **AMD64** folder and run **Setup.exe**.
2. Choose the language you want to use during installation. Click the down arrow and select the desired language from the drop-down list, and click **Next** to continue.



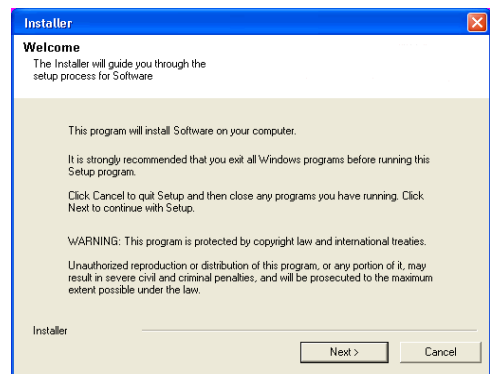
3. Select the option to install software on this computer.

**NOTES**


- The options that appear on this screen depend on the computer in which the software is being installed.



4. Read the Welcome screen.  
Click **Next** to continue, if no other applications are running.



5. Read the virus scanning software warning.  
Click **OK** to continue, if virus scanning software is disabled.



6. Read the license agreement, then select **I accept the terms in the license agreement**.  
Click **Next** to continue.




## CLUSTER SELECTION

If you are installing in clustered environment, follow the steps below. For non-clustered environment, skip to Select Components for Installation.


7. Select **Configure a Virtual Server**.

Click **Next** to continue.



8. Select the disk group in which the cluster group resides.

Click **Next** to continue.



## SELECT COMPONENTS FOR INSTALLATION

9. Select the component(s) to install.

### NOTES

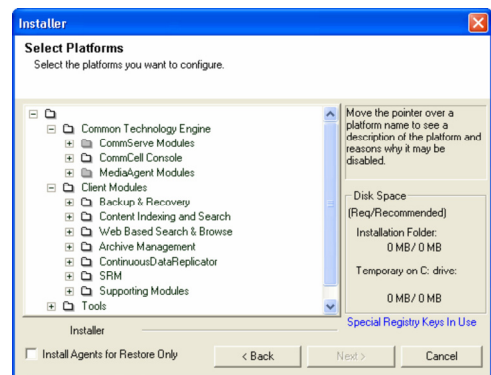
- Your screen may look different from the example shown.
- Components that either have already been installed, or which cannot be installed, will be dimmed. Hover over the component for additional details.
- If you wish to install the agent software for restore only, select **Install Agents for Restore Only** checkbox. See Installing Restore Only Agents for more information.
- The **Special Registry Keys In Use** field will be highlighted when GalaxyInstallerFlags registry key is enabled. Move the mouse pointer over this field to see a list of registry keys that have been created in this computer.

Click **Next** to continue.

To install the Quick Recovery Agent, expand the Client Modules folder and the Quick Recovery folder and select the following:

- Quick Recovery Agent

When you select the Quick Recovery Agent for installation, QSnap is automatically



selected for installation.

Select for installation any of the following components that you plan to use with the Quick Recovery Agent:

- SnapView Enabler for Quick Recovery
- VSS Enabler for Quick Recovery
- SnapVault/SnapMirror ONTAP Enabler
- SnapVault Open Systems Enabler

The SnapVault Open Systems Enabler and the NDMP Remote Server should not be installed on the same client. Only one of the enablers can be installed and can work on a given client.

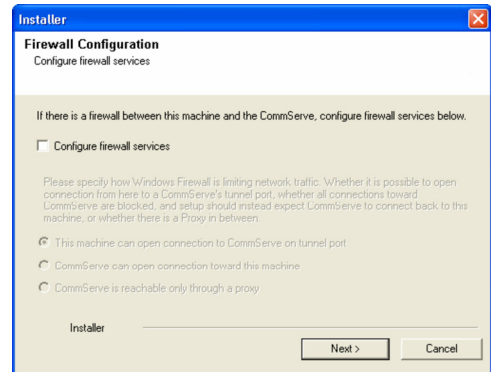
If you intend to use the Quick Recovery Agent with an application agent for application level browse and recover, see Installation for the installation procedure for that agent.

## CONFIGURATION OF OTHER INSTALLATION OPTIONS

10. If this computer and the CommServe is separated by a firewall, select the **Configure firewall services** option and then click **Next** to continue.

For firewall options and configuration instructions, see Firewall Configuration and continue with the installation.

If firewall configuration is not required, click **Next** to continue.

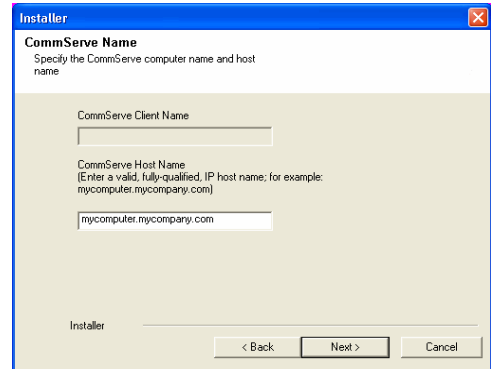


11. Enter the fully qualified domain name of the CommServe Host Name. This should be TCP/IP network name. e.g., computer.company.com.

### NOTES

- The CommServe client name is the name of the computer. This field is automatically populated.
- Do not use space and the following characters when specifying a new name for the CommServe Host Name:  
`|\`~!@#$$%^&*()+=<>/?,[\]{}:;'"`
- If a computer has already been installed, this screen will not be displayed; instead the installer will use the same Server Name as previously specified.
- If you do not specify the CommServe Host Name, a window will be prompted to continue in decouple mode. Click **Yes** to continue to Decoupled Install. Click **No** to specify a CommServe Name and continue with the installation.

Click **Next** to continue.

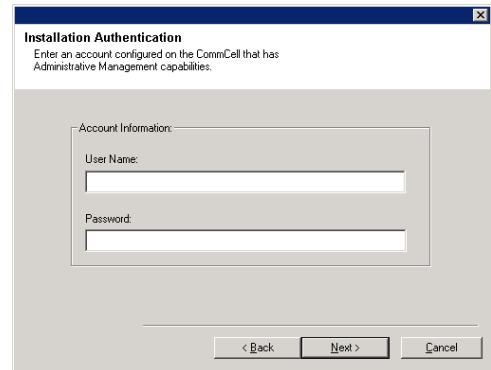


12. Enter the **username** and **password** associated with an external domain user account or a CommCell user account to authorize the installation of this agent.

### NOTES

- This window will be displayed when the **Require Authentication for Agent Installation** option is selected in the **CommCell Properties**. For more information, see Authentication for Agent Installs.

Click **Next** to continue.




13. Enter the following:

- The local (NetBIOS) name of the client computer.
- The TCP/IP IP host name of the NIC that the client computer must use to communicate with the CommServe Server.

### NOTES

- Do not use spaces when specifying a new name for the Client.
- The default network interface name of the client computer is displayed if the computer has only one network interface. If the computer has multiple network interfaces, enter the interface name that is preferred for communication with the CommServe Server.
- If a component has already been installed, this screen will not be displayed; instead, the install program will use the same name as previously specified.

Click **Next** to continue.



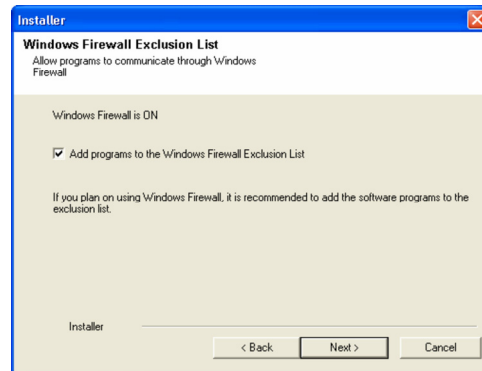
14. Select **Add programs to the Windows Firewall Exclusion List**, if you wish to add CommCell programs and services to the Windows Firewall Exclusion List.

**NOTES:**

- If Windows Firewall is enabled on the computer, this option is selected by default and must be enabled to proceed with the installation.
- If Windows Firewall is disabled on the computer, you can select this option to add the programs and services to enabled CommCell operations across the firewall, if the firewall is enabled at a later time.

You can either select this option during install or add the programs and services after installation. For adding the programs and services after installation, see Configure Windows Firewall to Allow CommCell Communication.

Click **Next** to continue.



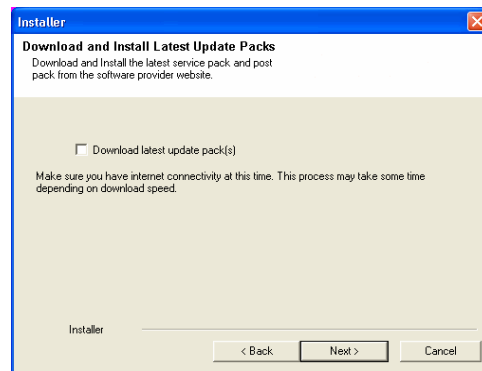
**DOWNLOAD AND INSTALL LATEST PACKS**

15. Select **Download latest update pack(s)** to automatically download and install the latest service packs and/or post packs if applicable at the end of this agent install.

**NOTES**

- Internet connectivity is required to download updates.
- Updates are downloaded to the following directory:  
<software installation>/Base/Temp/DownloadedPacks.  
They are launched silently and installed automatically for the first instance.

Click **Next** to continue.



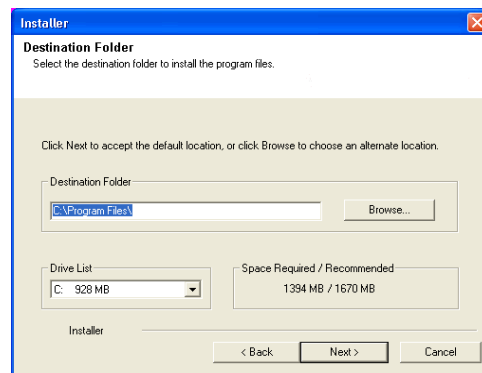
16. Specify the location where you want to install the software.

**NOTES**

- Do not install the software to a mapped network drive.
- Do not use the following characters when specifying the destination path:  
/ : \* ? " < > | #  
It is recommended that you use alphanumeric characters only.
- If you intend to install other components on this computer, the selected installation directory will be automatically used for that software as well.
- If a component is already installed in this computer, this screen may not be displayed. The software will be automatically installed in the same location that was previously specified.

Click **Browse** to change directories.

Click **Next** to continue.




**CLIENT GROUP SELECTION**

17. Select a Client Group from the list.

Click **Next** to continue.

**NOTES**

- This screen will be displayed if Client Groups are configured in the CommCell Console. For more information, see Client Computer Groups.




**SCHEDULE AUTOMATIC UPDATE**

18. If necessary, select this option to schedule an automatic installation of software updates.

**NOTES**

- Schedule Install of Automatic Updates allows automatic installation of the necessary software updates on the computer on a single or weekly basis. If you do not select this option, you can schedule these updates later from the CommCell Console.
- To avoid conflict, do not schedule the automatic installation of software updates to occur at the same time as the automatic FTP downloading of software updates.
- If a component has already been installed, this screen will not be displayed; instead, the installer will use the same option as previously specified.

Click **Next** to continue.



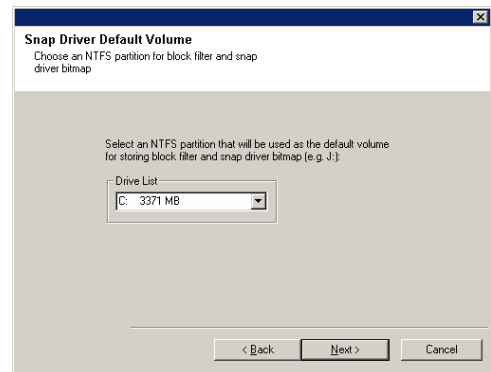
**CONFIGURE QSNAP**

19. From the drop-down list, select an NTFS partition that will be used as the default volume for bitmap file storage.

**NOTES**

- For standard installation, you can select any available NTFS formatted drive as the default volume for bitmap file storage.
- For cluster installation, the default location for storing the bitmap file is the corresponding shared volume. After the installation is complete, see Change the QSnap Bitmap Location for step-by-step instructions on changing the bitmap location.
- Only NTFS volumes will be shown in the drop-down list.

Click **Next** to continue.



**VERIFY SUMMARY OF INSTALL OPTIONS**

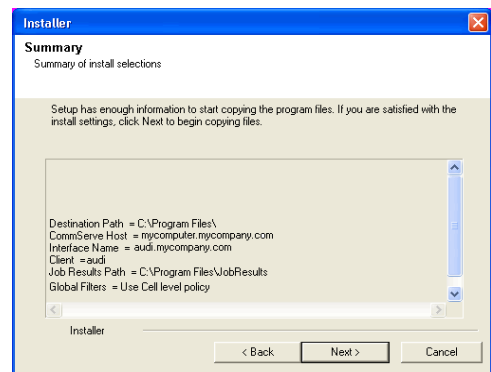
20. Verify the summary of selected options.

**NOTES**

- The **Summary** on your screen should reflect the components you selected for install, and may look different from the example shown.

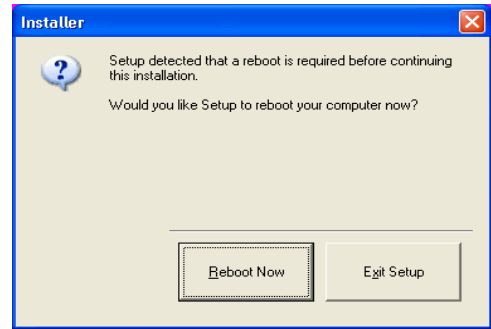
Click **Next** to continue or **Back** to change any of the options.

The install program now starts copying the software to the computer. This step may take several minutes to complete.



21. The System Reboot message may be displayed. If so, select one of the following:


- **Reboot Now**  
If this option is displayed without the **Skip Reboot** option, the install program has found files required by the software that are in use and need to be replaced. If **Reboot Now** is displayed without the **Skip Reboot** option, reboot the computer at this point. The install program will automatically continue after the reboot.
- **Exit Setup**  
If you want to exit the install program, click **Exit Setup**.



## INSTALL REMAINING CLUSTER NODES

If you are installing in clustered environment, follow the steps below to install on remaining nodes of the cluster. For non-clustered environment, skip to Setup Complete.

22. To install/upgrade the software on the remaining nodes of the cluster, click **Yes**.  
To complete the install for this node only, click **No**.




23. Select cluster nodes from the **Preferred Nodes** list and click the arrow button to move them to the **Selected Nodes** list.

### NOTES


- The list of **Preferred Nodes** displays all the nodes found in the cluster; from this list you should only select cluster nodes configured to host this cluster group server.
- Do not select nodes that already have multiple instances installed. For more information, see Multi Instancing.

When you have completed your selections, click **Next** to continue.



24. Type the **User Name** and **Password** for the Domain Administrator account, so that the installer can perform the remote install/upgrade of the cluster nodes you selected in the previous step.

Click **Next** to continue.



25. The progress of the remote install for the cluster nodes is displayed; the install can be interrupted if necessary.

Click **Stop** to prevent installation to any nodes after the current ones complete.

Click **Advanced Settings** to specify any of the following:

- Maximum number of nodes on which Setup can run simultaneously.
- Time allocated for Setup to begin executing on each node, after which the install attempt will fail.
- Time allocated for Setup to complete on each node, after which the install attempt will fail.

### NOTES

- If, during the remote install of a cluster node, setup fails to complete or is interrupted, you must perform a local install on that node. When you do, the install begins from where it left off, or from the beginning if necessary. For procedures,

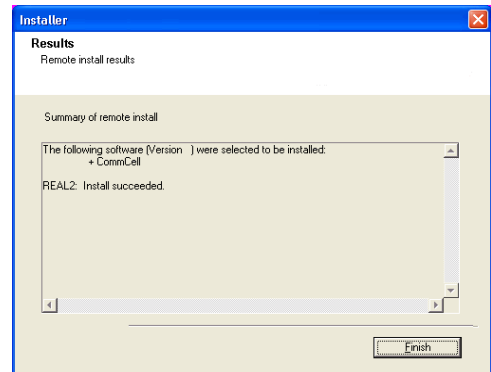


see Manually Installing the Software on a Passive Node.

26. Read the summary for remote installation to verify that all selected nodes were installed successfully.

**NOTES**

- If any node installation fails, you must manually install the software on that node once the current installation is complete. (See Manually Installing the Software on a Passive Node for step-by-step instructions.)
- The message displayed on your screen will reflect the status of the selected nodes, and may look different from the example.

Click **Next** to continue.



**SETUP COMPLETE**

27. Click **Next** to continue.

**NOTES**

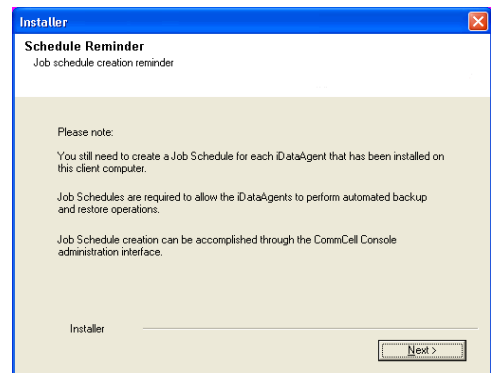
- Schedules help ensure that the data protection operations for the Agent are automatically performed on a regular basis without user intervention. For more information, see Scheduling.

28. Select from the following:

- If the **Reboot Now** button is displayed, a reboot is required before you can use the software. You can click this button to restart the computer now, or choose to perform the restart at another time. If the **Reboot Now** button is not displayed, it will not be necessary to restart the computer.
- Click **Finish** to exit the program.

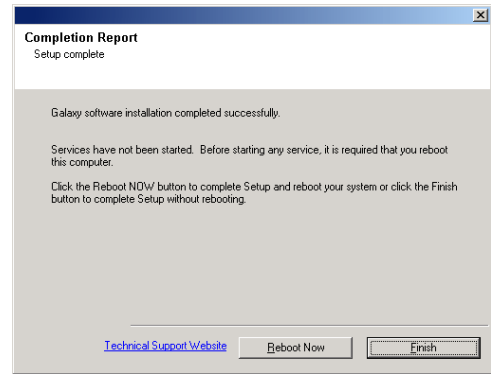
**NOTES**

- The **Setup Complete** message displayed on your screen will reflect the components you installed/upgraded, and may look different from the example shown.
- If you install an Agent with the CommCell Console open, you need to refresh the CommCell Console (F5) to see the new Agents.





This procedure is now complete.



## POST-INSTALL CONSIDERATIONS

---

### GENERAL

- Review Install Considerations after installing the software.
- Install post-release updates or Service Packs that may have been released after the release of the software. When you are installing a Service Pack, ensure that it is the same version as the one installed in the CommServe Server. Alternatively, you can enable Automatic Updates for quick and easy installation of updates in the CommCell component.

# Install the Quick Recovery Agent - Unix

## TABLE OF CONTENTS

### Install Requirements

#### Before You Begin

#### Install Procedure

- Getting Started
- Select Components for Installation
- Base Software Installation
- Kernel Parameters
- Enable Global Filters
- Client Group Selection
- Storage Policy Selection
- Setup Complete

#### Post-Install Considerations

## INSTALL REQUIREMENTS

The following procedure describes the steps for installing the Quick Recovery Agent. The Quick Recovery Agent is installed on the computer on which the volume to be protected is mounted. (This computer is referred to as the *Client* computer in this install procedure.)

Verify that the computer in which you wish to install the software satisfies the minimum requirements specified in System Requirements - Quick Recovery Agent.

Review the following before installing the software:

---

### GENERAL

- Review Install Considerations before installing the software.
- Agents should be installed only after the CommServe and at least one MediaAgent have been installed in the CommCell. Also, keep in mind that the CommServe and MediaAgent must be installed and running (but not necessarily on the same computer), before you can install the Agent.
- Ensure there is an available license on the CommServe for the Agent.
- Verify that you have the Software Installation Disc that is appropriate to the destination computer's operating system.

---

### HP-UX

- If you are installing on a HP-UX computer, you must manually mount the installation disc as described in Mount the Software Installation Disc.

## BEFORE YOU BEGIN

- Log on to the client as `root`.
- The install package requires `super-user` permissions to execute.

## INSTALL PROCEDURE

---

### GETTING STARTED

1. Place the software installation disc for the Unix platform into the disc drive.  
You can also install the product using a disc drive mounted on another computer on the network.
  - On Solaris, double-click the **cvpkgadd** program from the File Manager window.
  - On other Unix platforms, open the Terminal window, navigate to the software installation disc and then enter **./cvpkgadd**.
2. The product banner and other information is displayed.  
Press **Enter** to continue.
3. Read the license agreement. Type **y** and press **Enter** to continue.
4. Enter the number corresponding to the setup task you want to perform.

Please select a setup task you want to perform from the list below:

Advance options provide extra setup features such as creating custom package, recording/replaying user selections and installing External Data Connector software.

- 1) Install data protection agents on this computer
- 2) Advance options

### NOTES

- For Install data protection agents on this computer option, follow the steps described in this procedure.
- Advance options provide additional setup features such as record and play setup, creating a custom package and External Data Connector Agent software.

To create a custom package and for record and play setup, follow the steps

described in Custom Package - Unix.

To install the External Data Connector Agent, follow the steps described in External Data Connector - Unix.

- If your computer is 32-bit, press **Enter**.

If your computer is 64-bit, see Install Unix Agents on 64-bit Platform for step-by-step procedure.

- This prompt is displayed only when you are installing on AIX, HP-UX, Linux, or Solaris computers.

Press **Enter** to continue

**NOTES**

- When you install on non-clustered computer, you must select the number associated with the option **Install on a physical machine**.

- If you have only one network interface, press **Enter** to accept the default network interface name and continue.

If you have multiple network interfaces, enter the number corresponding to the network interface that you wish to use as default, and then press **Enter** to continue.

**NOTES**

- The interface name and IP addresses depend on the computer in which the software is installed and may be different from the example shown.

- Specify the client name for the computer.

Press **Enter** to accept the default name and continue, or Enter a new client name for the computer and then press **Enter** to continue.

3) Exit this menu

Your choice: [1]

This machine supports both 32 bit and 64 bit binaries. By default, we will install 32 bit binary set that has full support for all the modules included in this package. Please note that 64 bit binary set currently only support limited modules.

1) All platforms (32 bit)

2) FS and MA only (64 bit)

Your choice: [1]

Certain Calypso packages can be associated with a virtual IP, or in other words, installed on a "virtual machine" belonging to some cluster. At any given time the virtual machine's services and IP address are active on only one of the cluster's servers. The virtual machine can "fail-over" from one server to another, which includes stopping services and deactivating IP address on the first server and activating the IP address/services on the other server.

You now have a choice of performing a regular Calypso install on the physical host or installing Calypso on a virtual machine for operation within a cluster.

Most users should select "Install on a physical machine" here.

1) Install on a physical machine

2) Install on a virtual machine

3) Exit

Your choice: [1]

We found one network interface available on your machine. We will associate it with the physical machine being installed, and it will also be used by the CommServe to connect to the physical machine. Note that you will be able to additionally customize Datapipe Interface Pairs used for the backup data traffic later in the Calypso Java GUI.

Please check the interface name below, and make connections if necessary:

Physical Machine Host Name: [angel.company.com]

Please specify the client name for this machine.

It does not have to be the network host name: you can enter any word here without spaces. The only requirement is that it must be unique on the CommServe.

Physical Machine Client name: [angel]

**SELECT COMPONENTS FOR INSTALLATION**

- Enter the number corresponding to the **CVGxQRA** module.

A confirmation screen will mark your choice with an "X". Type "d" for **Done**, and press **Enter** to continue.

**NOTES**

- To select multiple component, enter the number by adding a space.
- Your screen may look different from the example shown.
- Components that either have already been installed, or which cannot be installed, will not be shown.
- In addition, the list of modules that appear depends on the specific Unix File System in which the package is installed. (e.g., **CVGxWA** will appear only when the installation package is run on a Solaris computer.)

Install Calypso on physical machine client.company.com

Select the Calypso module that you would like to install

```
[ ] 1) Media Agent          [1301] [CVGxMA]
[ ] 2) FileSystem IDA      [1101] [CVGxIDA]
>) >>>> NEXT PAGE >>>>>
```

[a=all n=none r=reverse q=quit d=done >=next <=previous ? =help]

Enter number(s)/one of "a,n,r,q,d,>,<,>?" here: 2

**BASE SOFTWARE INSTALLATION**

- If you wish to install the agent software for restore only, enter **Yes** and press **Enter** to continue. See Installing Restore Only Agents for more information.

Otherwise, accept **no**, press **Enter** to continue.

- Type the appropriate number to install the latest software scripts and press **Enter** to continue.

**NOTES**

- Select **Download from the software provider website** to download the latest

Do you want to use the agents for restore only without consuming licenses? [no]

Installation Scripts Pack provides extra functions and latest support and fix performed during setup time. Please specify how you want to get this pack.

If you choose to download it from the website now, please make sure you have internet connectivity at this time.

software scripts from your software provider website.

Make sure you have internet connectivity when you are using this option.

- Select **Use the one in the installation media**, to install the software scripts from the disc or share from which the installation is performed.
- Select **Use the copy I already have by entering its unix path**, to specify the path if you have the software script in an alternate location.

12. Enter **Yes** to download and install the latest service packs and post packs from the software provider.

**NOTES**

- Internet connectivity is required to download updates.
- This step is applicable for multi instancing.

Press **Enter** to continue.

13. Specify the location where you want to install the software.

**NOTES**

- The amount of free space required depends on the components selected for install, and may look different from the example shown.

Press **Enter** to accept the default path and continue, or Enter a path and then press **Enter** to continue.

Press **Enter** again to confirm the path.

14. Specify the location for the log files.

**NOTES**

- All the modules installed on the computer will store the log files in this directory.
- The amount of free space required depends on the components selected for install, and may look different from the example shown.

Press **Enter** to accept the default path and continue, or Enter a path and then press **Enter** to continue.

Press **Enter** again to confirm the path.

15. Indicate whether you would like to launch processes with inherent database access rights.

Press **Enter** to assign a new group, or Type **No** and then press **Enter** to continue.

16. If you indicated **Yes** in the previous step, you will be prompted for the group name that must be used to launch processes.

Enter the group name and then press **Enter** to continue.

Press **Enter** again to continue.

For installs on a Solaris computer, proceed to the next step. Otherwise, skip to Storage Policy Selection.

17. Type a network TCP port number for the Communications Service (CVD) and press **Enter**.

Type a network TCP port number for the Client Event Manager Service (EvMgrC) and press **Enter**.

This process may take some time depending on the internet connectivity.

- 1) Download from the software provider website.
  - 2) Use the one in the installation media
  - 3) Use the copy I already have by entering its unix path
- Your choice: [1] 2

Keep Your Install Up to Date - Latest Service Pack

Latest Service Pack provides extra functions and latest support and fix for the packages you are going to install. You can download the latest service pack from software provider website.

If you decide to download it from the website now, please make sure you have internet connectivity at this time. This process may take some time depending on the internet connectivity.

Do you want to download the latest service pack now? [no]  
Press <ENTER> to continue ...

Please specify where you want us to install Calypso binaries.

It must be a local directory and there should be at least 98MB of free space available. All files will be installed in a "calypso" subdirectory, so if you enter "/opt", the files will actually be placed into "/opt/calypso".

Installation Directory: [/opt]

..

Calypso will be installed in /opt/calypso.  
Press ENTER to continue ...

Please specify where you want to keep Calypso log files.

It must be a local directory and there should be at least 100MB of free space available. All log files will be created in a "calypso/Log\_Files" subdirectory, so if you enter "/var/log", the logs will actually be placed into "/var/log/calypso/Log\_Files".

Log Directory: [/var/log]

..

Calypso log files will be created in /var/log/calypso/Log\_Files.  
Press ENTER to continue ...

Most of Calypso processes run with root privileges, but some are launched by databases and inherit database access rights. To make sure that registry and log files can be written to by both kinds of processes we can either make such files world-writeable or we can grant write access only to processes belonging to a particular group, e.g. a "calypso" or a "dba" group.

We highly recommend now that you create a new user group and enter its name in the next setup screen. If you choose not to assign a dedicated group to Calypso processes, all temporary and configuration files will be created with -rw-rw-rw permissions.

If you're planning to backup Oracle DB you should use "dba" group.

Would you like to assign a specific group to Calypso? [yes]

Please enter the name of the group which will be assigned to all Calypso files and on behalf of which all Calypso processes will run.

In most of the cases it's a good idea to create a dedicated "calypso" group. However, if you're planning to use Oracle iDataAgent or SAP Agent, you should enter Oracle's "dba" group here.

Group name: dba

REMINDER

If you are planning to install Calypso Informix, DB2, PostgreSQL, Sybase or Lotus Notes iDataAgent, please make sure to include Informix, DB2, etc. users into group "dba".

Press <ENTER> to continue ...

Every instance of Calypso should use a unique set of network ports to avoid interfering with other instances running on the same machine. The port numbers selected must be from the reserved port number range and have not been registered by another application on this machine.

**NOTES**

- For more information about Network TCP Ports, see Network TCP Port Requirements.
- For more information about these services, see Services.
- If the port number you entered already exists, a message will be displayed `Port ### is already reserved in /etc/services`. To work around this issue, enter different port number.

18. If this computer and the CommServe is separated by a firewall, type **Yes** and then press **Enter** to continue.

For firewall options and configuration instructions, see Firewall Configuration and continue with the installation.

If you do not wish to configure the firewall services, type **No** and then press **Enter** to continue.

19. Type the name of the CommServe computer and press **Enter** to continue.

**NOTES**

- Ensure that the CommServe is accessible before typing the name; otherwise the installation will fail.
- If you enter a short name which resolves to the same IP address as the fully qualified CommServe name, you will be asked if you would prefer to use the fully qualified name.

20. Enter the **username** and **password** information for an external domain user account or a CommCell user account. This authorizes the installation of an agent on the CommCell.

**NOTES**

- This is only displayed when the **Authentication for Agent** feature is enabled in the CommCell Properties. Users must belong to a User Group with Agent Management capabilities to enable this feature. For more information, see Authentication for Agent Installs.

Click **Enter** to continue.

Please enter the port numbers.

Port Number for CVD : [8600]

Port Number for EvMgrC: [8602]

Is there a firewall between this client and the CommServe? [no]

Please specify hostname of the CommServe below. Make sure the hostname is fully qualified, resolvable by the name services configured on this machine.

CommServe Host Name:

Enter your CommCell user name and password:

User Name :

Password :

Press <ENTER> to continue ...

**KERNEL PARAMETERS**

21. Enter the appropriate number of streams, and then press **Enter** to continue, or Press **Enter** to accept the default number of streams and continue.

**NOTES**

- The number of streams specified ensures that concurrent backup/restore streams would have sufficient system resources. For more information on the subject, see Configuring Kernel Parameters for Macintosh and Configuring Kernel Parameters for Solaris.

This prompt is relevant only when you install/upgrade on a Macintosh or Solaris computer as appropriate.

22. Indicate whether you would like modifications to be made to the `/etc/system` configuration file.

Type **Yes**, and then press **Enter** to automatically update the file and continue, or Press **Enter** to accept the default **No** and continue (if you do not want to automatically update the file).

This prompt is displayed only when you install/upgrade on a Solaris (8 or 9) or Macintosh computer.

Please enter the total number of streams that you plan to run at the same time. We need to make sure that you have enough semaphores and shared memory segments configured in `/etc/system`.

Number of streams: [10]

We now need to modify the `/etc/system` configuration file on this computer. It is done to make sure that there will be enough shared memory and semaphores available for Calypso programs.

Please review the changes below and answer "yes" if you want us to apply them to the `/etc/system` file. Otherwise, the installation will proceed, the changes will be saved to some other file, and you will have to apply them manually.

```
set shmsys:shminfo_shmmni=8570 (was 7930)
set shmsys:shminfo_shmseg=8420 (was 7780)
set semsys:seminfo_semmns=10320 (was 9680)
set semsys:seminfo_semmni=8570 (was 7930)
set semsys:seminfo_semmns1=8570 (was 7930)
```

Do you want us to apply these changes now? [no]

Changes saved into `/etc/system.gal.1744`

Press <ENTER> to continue.

23. If you indicated **No** in the previous step, the file to which the changes have been saved is displayed. Make sure that these values are established later to ensure that all the requirements for this setup is satisfied.

**NOTES**

- The settings that are displayed are the maximum or minimum required settings. Value '640', which is provided for various shared memory segment or semaphore requirements, is a maximum value based on 10 streams.

Although a 'no' answer can be selected to this question during install, the user should make sure the min requirements (below) for shared memory are met, otherwise the backups may fail (the message in logs is 'could not start the pipeline').

```
set shmsys:shminfo_shmmax=4199304
set shmsys:shminfo_shmmni=1
set semsys:shminfo_shmmni=640
set semsys:shminfo_shmseg=640
set semsys:seminfo_semmns=640
set semsys:seminfo_semmni=640
```

Press **Enter** to continue.

This prompt is displayed only when you install/upgrade on a Solaris (8 or 9) computer, in cases where the install detects that the computer does not have the maximum or minimum required shared memory settings.

```
set semsys:seminfo_semmsl=640
set maxusers=256
```

## ENABLE GLOBAL FILTERS

24. Type the appropriate number for configuring Global Filters for the default subclient and press Enter to continue.

### NOTES

- Select **Use Cell level Policy** to inherit the global filter policy configuration set for the CommCell, i.e., if the **Use Global Filters on All Subclients** option is selected in the **Global Filters** dialog box (from the CommCell Console's Control Panel), then this policy will be applied to the default subclient as well. If is not selected, then the global filters will not be applied to the default subclient.
- Select **Always use Global filters** to always apply the global filters policy to the default subclient regardless of the policy set for the CommCell.
- Select **Do not use Global filters** to disregard applying the global filters to the default subclient regardless of the policy set for the CommCell.

Commcell Level Global Filters are set through Calypso GUI's Control Panel in order to filter out certain directories or files from backup Commcell-widely. If you turn on the Global filters, they will be effective to the default subclient. There are three options you can choose to set the filters.

- 1) Use Cell level policy
- 2) Always use Global filters
- 3) Do not use Global filters

Please select how to set the Global Filters for the default subclient? [1]

## CLIENT GROUP SELECTION

25. Type the number of a Client Group and press **Enter**.

A confirmation screen will mark your choice with an "X". Type **d** for done with the selection, and press **Enter** to continue.

### NOTES

- This screen will be displayed only if Client Groups are configured for the CommCell. For more information, see Client Computer Groups.

Client Group(s) is currently configured on CommServe cs.company.com. Please choose the group(s) that you want to add this client client.company.com to. The selected group(s) will be marked (X) and can be deselected if you enter the same number again. After you are finished with the selection, select "Done with the Selection".

- ```
[ ] 1) Unix
[ ] 2) DR
[ ] 3) DKS
```

```
[a=all n=none r=reverse q=quit d=done >=next <=previous ?
=help]
```

Enter number(s)/one of "a,n,r,q,d,>,<," here: 2

```
+-----+
```

IMPORTANT:

In addition to installing Calypso on this computer, you will also need to create a Job Schedule for each iDataAgent that has been installed on this client computer.

Job Schedules are required to allow the Calypso iDataAgents to perform automated backup and restore operations.

Job Schedule creation can be accomplished through the Calypso CommCell Console administration interface.

```
+-----+
```

26. Press **Enter** to continue.

### NOTES

- Schedules help ensure that the data protection operations for the Agent are automatically performed on a regular basis without user intervention. For more information, see Scheduling.

## STORAGE POLICY SELECTION

27. Enter the number corresponding to the storage policy through which you want to back up the File System iDataAgent and then press **Enter** to continue.

### NOTES

- A storage policy directs backup data to a media library. Each library has a default storage policy.
- When you install an Agent, the install program creates a default subclient for most Agents.
- If desired, you can change your storage policy selection at any time after you have installed the client software.
- If this screen appears more than once, it is because you have selected multiple agents for installation and are configuring storage policy association for each of the installed agents.

Please select one storage policy for this IDA from the list below:

- 1) SP\_StandAloneLibrary2\_2
- 2) SP\_Library3\_3
- 3) SP\_MagLibrary4\_4
- 4) fornax\_fornax\_HWCmp
- 5) ranger\_ranger\_HWCmp
- 6) fornax\_fornax\_CIntCmp
- 7) fornax\_fornax\_MACmp
- 8) fornax\_fornax\_NoCmp

Storage Policy: [3]

## SETUP COMPLETE

28. Installation of the QR Agent is now complete. You are warned to install the Oracle iDataAgent if you intend to run Oracle log backups.

Press **Enter** to continue.

```
Updating registry tree under /etc/CommVaultRegistry ...
done.
Preconfiguring CVGxQRA on the CommServe ... done.
```

WARNING!

If you plan to run Oracle log backups along with QRA incremental updates, you will have to install Calypso Oracle iDataAgent (CVGxOrIDA) to perform successful restores.

If this is the case, we suggest that you install CVGxOrIDA now.

Adjusting modes and permissions of Calypso files

Successfully installed CVGxQRA.

Press <ENTER> to proceed

Select the Calypso module that you would like to install.

- 1) FileSystem iDataAgent
- 2) Exit

Module number: [1]

29. This menu may be displayed only when you are installing on HP-UX, Linux, or Solaris computers. If this is the last package that you wish to install/upgrade, enter the number corresponding to the **Exit** option and then press **Enter** to continue.

**NOTES**

- Only modules that are not installed/upgraded appear in the list.
- Your screen may appear different from the example shown.
- If you are installing on AIX, FreeBSD, IRIX or Tru64 computers, if this module was the last possible module to install, you are automatically exited from the program. Otherwise, type the number for the **Exit** option and then press **Enter**. The installation is completed.

30. This prompt is displayed only when you are installing on HP-UX, Linux, or Solaris computers. Enter the number corresponding to the **Exit** option and then press **Enter** to continue.

The installation is now complete.

Certain Calypso packages can be associated with a virtual IP, or in other words, installed on a "virtual machine" belonging to some cluster. At any given time the virtual machine's services and IP address are active on only one of the cluster's servers. The virtual machine can "fail-over" from one server to another, which includes stopping services and deactivating IP address on the first server and activating the IP address/services on the other server.

Currently you have Calypso installed on physical node stone.company.com.

Now you have a choice of either adding another package to the existing installation or configure Calypso on a virtual machine for use in a cluster.

- 1) Add another package to stone.company.com
- 2) Install Calypso on a virtual machine
- 3) Exit

Your choice: [1]

## POST-INSTALL CONSIDERATIONS

### GENERAL

- Review Install Considerations after installing the software.
- Install post-release updates or Service Packs that may have been released after the release of the software. When you are installing a Service Pack, ensure that it is the same version as the one installed in the CommServe Server. Alternatively, you can enable Automatic Updates for quick and easy installation of updates in the CommCell component.

# Recovery Points

Topics | How To | Related Topics

This feature/product/platform is on Extended Support in this release. See [Deprecated Features, Products, and Platforms](#) for more information.

## Overview

### Recovery Points for the Quick Recovery Agent

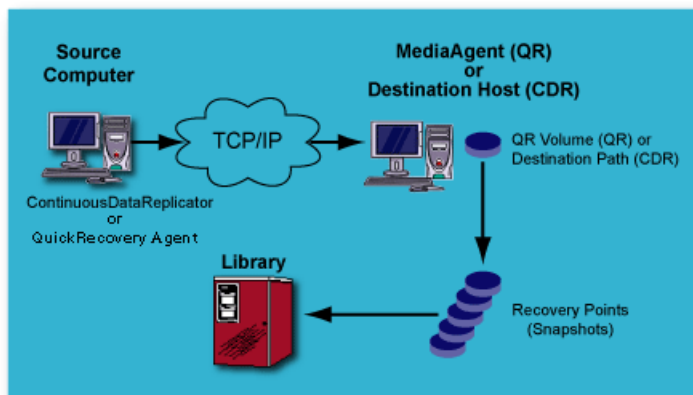
- Software and Hardware Requirements
- Implement Recovery Points Using the Quick Recovery Agent
- How to create Consistent Recovery Points using Quick Recovery Agent
- How to browse Quick Recovery Agent Consistent Recovery Points
- How to delete Quick Recovery Agent Consistent Recovery Points
- Recovery
  - Recovering Consistent Recovery Points Using Quick Recovery Agent
  - Other Quick Recovery Agent Recovery Options
- Best Practices

### License Requirement

## OVERVIEW

Recovery Points is a snapshot of the data to preserve a point-in-time on the QR Volume. This not only affords an extra measure of protection for your data, it also expands the number of options you have when recovering your data.

This illustration provides a high-level look at how data flows when creating Recovery Points with Quick Recovery Agent:



## RECOVERY POINTS FOR THE QUICK RECOVERY AGENT

Using QR Agent with QSnap, you can create Consistent Recovery Points to preserve the states of QR Volumes after incremental updates, providing more protection and flexibility for recovery. (See [Overview - QSnap](#) for more information about QSnap.) Without Consistent Recovery Points, recovery is limited to whatever is on your QR Volume after the last update. In addition, your QR Volume is no longer up to date with the source volume if there is a write to the QR Volume between update operations.

With the Recovery Points feature enabled in a QR Policy and with volumes updated incrementally, snapshots of the QR Volume are taken after the first full copy to it from the source and after each succeeding incremental update -- up to the number of Points you specify (maximum: 32). Each snapshot represents a Consistent Recovery Point because it is taken on the QR Volume after incremental updates are added to the volume.

Showing up as snapshots in the Browse QR Volumes window, Consistent Recovery Points can be mounted, shared on a network, or copied back. In addition, Consistent Recovery Points created with QR Agent **persist** through reboots. See [Mount and Unmount QR Volumes](#) and [Copy Back a QR Volume](#) for procedures relating to these topics.



---

## SOFTWARE AND HARDWARE REQUIREMENTS FOR RECOVER POINTS

To use the Recovery Point feature, the following is required:

- Source and QR Volume computers must meet the System Requirements for the Quick Recovery Agent. See System Requirements - Quick Recovery Agent.
- The Quick Recovery Agent and QSnap must be installed on both source and QR volume computers. (VSS may also be installed on the source volume computer.) See Install the Quick Recovery Agent - Windows and Install QSnap - Windows.
- Use of Copyback to recover a Volume is subject to the considerations for that QR Agent feature. See Copyback on the Recover QR Volumes page.
- A separate license is required for Recovery Points.
- Recovery Points are enabled in the QR Policy, and the number of Recovery Points, up to 32, is specified there.
- A separate volume must be specified for the destination client QSnap cache. See QSnap Copy-On-Write Cache on the Overview - QSnap page for more information on the QSnap cache.
- Recovery Points are only generated for QR Agent incremental updates. See QR Volume with Incremental Updates on the Create a QR Volume page for more information.

---

## IMPLEMENT RECOVERY POINTS USING THE QUICK RECOVERY AGENT

**Experienced QR Agent users:** Note the following when implementing Recovery Points with QR Agent:

- A separate license is required for Recovery Points.
- Recovery Points are enabled in the QR Policy, and the number of Recovery Points, up to 32, is specified there.
- A separate volume must be specified for the destination client QSnap cache. See QSnap Copy-On-Write Cache on the Overview - QSnap page for more information on the QSnap cache.
- Recovery Points are only generated for QR Agent incremental updates. See QR Volume with Incremental Updates on the Create a QR Volume page for more information.

---

## HOW TO CREATE CONSISTENT RECOVERY POINTS USING QUICK RECOVERY AGENT

Perform the following tasks to create Consistent Recovery Points using QR Agent:

1. Ensure that you have a Feature License for Recovery Points.
2. Use Volume Explorer to detect available volumes, including any new volumes you created.
3. Create a Scratch Volume Pool.
4. Create a QR Policy. Select the **Enable Recovery Points** check box, and type the number of Points (up to a maximum of 32) in the Number of Recovery Points box. (See screen example below these steps.) Also, the Copy Manager must be on the same computer as the QR Volume.
5. When you click OK in the QR Policy window after enabling Recovery Points, you will be prompted to enter a cache partition for snapshot operations (unless you had set this up when setting up the QR Volume). Note that the cache cannot be on the QR Volume because a volume that contains a live snapshot or live application cannot be locked, and the QR Volume must be locked when writing to it. See the sample QR Policy screen below.
6. Create a New Subclient.
7. Create a QR Volume That is Incrementally Updated.



One of the available Advanced Options when you create a QR Volume that is incrementally updated (and when the associated QR Policy has Recovery Points enabled) is the option to mount the latest Recovery Point to a specified mount point. This option, Mount Destination Snap, allows access to the destination snap during the next incremental QR Volume creation job when the destination volume is normally not accessible. This provides uninterrupted access to the destination even during the copy phase of the next incremental QR Volume creation job.

Successful completion of the above tasks causes the system to take and store snapshots of the QR Volume after the first full copy and after each incremental update — up to the number you specified (maximum: 32). These snapshots represent Consistent Recovery Points for your source data.

## HOW TO BROWSE QUICK RECOVERY AGENT CONSISTENT RECOVERY POINTS

Consistent Recovery Points are viewed in the Browse QR Volumes window on the source (production volume), even though the snapshots for the Consistent Recovery Points were taken on the QR Volume. You can right-click a specific Consistent Recovery Point, and choose Details to bring up the Volume Configuration window, which has General, Mount Points, and QR Volume tabs.



If you mount a Recovery Point from the Browse QR Volumes window, you can view its contents in Volume Explorer.

## HOW TO DELETE QUICK RECOVERY AGENT CONSISTENT RECOVERY POINTS

Consistent Recovery Points created with QR Agent are automatically deleted by the system in a FIFO order when the number of points you specified (or the maximum of 32) is reached. For example: if you specify 20 Recovery Points in the QR Policy window, when the system creates the 21st, the first Recovery Point that was created is deleted. Creation of the 22nd Recovery Point causes the second that was created to be deleted, and so on.

QR Agent Recovery Points can also be deleted manually — individually or all at once for a particular QR Volume.

- To delete Recovery Points one at a time, you must first unmount them. Then you can select and delete Points in the Browse QR Volumes window, but you must delete them in FIFO order. That is, you must first delete the oldest Point, then the next oldest, etc.
- To delete all the Recovery Points associated with a volume, either release the associated volume from within the Scratch Volume Pool, or unmount and delete the QR Volume in the Browse QR Volumes window.



Mounted Recovery Points cannot be deleted. If you try to delete a mounted Point, an error message appears.

## RECOVERY

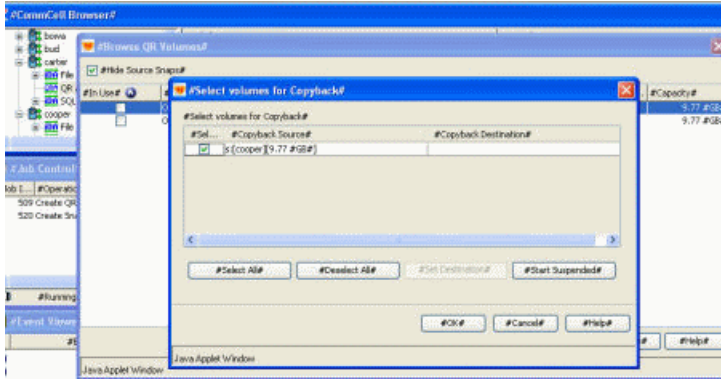
### RECOVERING CONSISTENT RECOVERY POINTS USING QUICK RECOVERY AGENT

**Experienced QR Agent users:** Note the following when recovering data from Recovery Points using QR Agent:

- In the Browse QR Volumes window, when you right-click a Recovery Point and select Recovery, the only option is Copyback; the Recover option is grayed out. See Copy Back a QR Volume.

Consistent Recovery Points created with the QR Agent are recovered in the Browse QR Volumes window through the Copyback feature; follow the Copy Back a

QR Volume procedure. Using Copyback, you can recover both mounted and unmounted Recovery Points.



It is possible to copy back to the source volume.

### OTHER QUICK RECOVERY AGENT RECOVERY OPTIONS

You can also recover data by manually copying and pasting data from mounted Recovery Points to other volumes.



Volume Explorer filters out QR Agent Recovery Points. But if a QR Agent Recovery Point is mounted to a drive letter, a detect operation in Volume Explorer will change the Recovery Point's associated QR volume entry so that both the QR volume's drive letter and the drive letter of the mount point are displayed. For example, if a QR volume is mounted to the J: drive and its associated Recovery Point is mounted to the P: drive, re-detecting volumes in Volume Explorer will show J:;P:; in the mount column.

---

### BEST PRACTICES

Keep in mind the following considerations when using the Recovery Points feature with the Quick Recovery Agent:

- Consistent Recovery Points are based on the incremental updates you set; you cannot create one on demand.
- You cannot specify that snapshots be taken only after specific incremental update operations. They are taken after every incremental update (and after the first full copy from source to QR Volume) up to the number you specified or the maximum of 32.
- The higher the number of snapshots specified, the more disk space you will need on the QR Volume's alternate cache partition.
- For information concerning QSnap's COW Cache, with reference to Recovery Points, refer to Considerations - QSnap COW Cache for Windows.

[Back to Top](#)

---

### LICENSE REQUIREMENTS

This feature requires a Feature License to be available in the CommServe® Server.

Review general license requirements included in License Administration. Also, View All Licenses provides step-by-step instructions on how to view the license information.

[Back to Top](#)

---

## Recovery Points for QR - How To

Topics | [How To](#)

---

[Create a QR Volume of a Subclient](#)

[Create a QR Policy](#)

[Create a New Subclient](#)

[Modify QR Policy Properties](#)

[Schedule QR Volume Creation](#)

### CREATE A QR VOLUME OF A SUBCLIENT

**Before you Begin:**

- Make sure you check QR Volume Creation Considerations for your agent.

Create QR Volume(s) allows you to secure data immediately without having to wait for the scheduled QR Volume Creation time. This capability can be useful if:

- You have some particularly valuable volumes that you need to secure immediately.
- You want to create a QR Volume that is not routinely secured through a QR Volume schedule.

*Required Capability:* See Capabilities and Permitted Actions

▶ To create a QR Volume of a selected subclient:

1. From the CommCell Browser, right-click the subclient that you want to copy and then click **Create QR Volume(s)**.
2. From the QR Volume Creation window, select when you want to run the job by clicking **Run Immediately**.
3. Select **Create QR Volume(s)**.
4. Optionally, click **Advanced** to access the QR Volume Creation Advanced Options dialog box.
5. After selecting the QR Volume type and any advanced options, click **OK**. The QR Volume operation will begin and you can track the progress of the job in the Job Controller window.

**CREATE A QR POLICY**

*Required Capability:* Capabilities and Permitted Actions

The system allows you to create a QR policy in the CommCell Console from the QR Policies level.

▶ To create a QR policy:

1. From the CommCell Browser, right-click the **QR Policies** icon, and then click **Create New QR Policy** from the short-cut menu.
2. Type the policy name (Up to 32 characters).
3. Select the **Snapshot Engine Type** from the dropdown menu.



- If you are creating snapshot scripts using the Generic Enabler feature, select **Generic snapshot** as your Snapshot Engine Type.
- Depending on the Snapshot Engine you select, some options may not be available for configuration.

4. Select the **Retention Policy** by checking Infinite (default), or by entering the number of hours/days/weeks.
5. Select the **Enable QR Volume Creation** check box to activate the **Copy Manager** and **Scratch Pool** menus.



This option is automatically disabled if you selected **Generic snapshot** as your Snapshot Engine Type in Step 3.

6. Select the **Copy Manager** from the list of available Copy Managers that currently exist for this QR Volume. (Refer to LAN Copy Manager.)
7. Select the **Scratch Pool** from the list of available scratch pools that currently exist for this QR Volume.
8. If this QR Policy will be used to create Recovery Points, select **Enable Recovery Points**, and specify the maximum number of Recovery Points that will be retained, up to a maximum of 32.
  - When you select **Enable Recovery Points**, you are prompted to enter a location for the cache partition for the snapshots of the QR Volume. This location cannot be on the QR Volume.
  - When you are satisfied with the QR Policy's configuration, click **OK**. The new QR Policy appears in the CommCell Browser.

**CREATE A NEW SUBCLIENT****Before You Begin**

- Review Subclients.
- Do not create a subclient while the parent node or any sibling subclient has a data protection or archive operation currently running on it.
- In cases where a new subclient is created with the same name as a deleted subclient, the system will append a Unix time stamp to the deleted subclient's name in data protection job history reports and views to distinguish the two subclients. For example, `subclientname_1104257351`.
- Informix *iDataAgents*: If you will be using the Informix ONBAR utility to create backup and restore scripts, you need not create subclients. Otherwise, if you will be using the CommCell Console to back up and restore Informix database objects (subsets/dbspaces), then you will need to create a subclient.
- ProxyHost *iDataAgents*: If you are using a BCV, you must prepare a batch file or a shell script file on the backup host containing commands to synchronize

and split the BCV. The Resource Pack includes information on configurations for these batch files or shell scripts, as well as examples that apply to specific applications and hardware (e.g., Exchange databases in an EMC Symmetrix environment). See Resource Pack for more information on the Resource Pack.

The ProxyHost *iDataAgent* also requires that you set permissions for the batch/shell script file on the backup host.

- SQL Server Database *iDataAgents*: When running on Windows Server 2003 and VSS is enabled, the **New Subclient** command is not available.
- PostgreSQL *iDataAgents*: Once you configure the PostgreSQL instance, the system automatically generates the default backup sets and default subclients. However, you can use the CommCell Console to create user-defined subclients for dump backup sets to distribute some of the database content. You cannot create user-defined subclients for FS backup sets.

*Required Capability:* See Capabilities and Permitted Actions

▶ To create a new subclient:

1. From the CommCell Browser, right-click the node (agent/backup set/archive set/instance) for which you want to create a new subclient, click **All Tasks** (if applicable), and then simply click **New Subclient** for most agents.
  - For the SQL Server *iDataAgent*, expand **New Subclient** and click either **Database** to include individual databases or **File/File Group** to include database elements.
2. Click the General tab or General (Quick Recovery Agent) tab of the Subclient Properties dialog box and type the name (up to 32 characters) of the subclient that you want to create.
  - For supported agents identified in Support Information - Snapshot Engines, you can select a QSnap option to snap data and then perform a data protection operation on the data.
  - For Image Level on Unix and Image Level ProxyHost on Unix, use the **Incremental Support Using** field to configure either a CXBF subclient or a checksum subclient and to enable incremental support for either subclient type.
  - For QR Agents, you must also select a QR Policy from the **QR Policy** list.
  - For the Windows *iDataAgents* that support VSS, you can optionally Enable VSS on a Subclient.

3. Select other options from the General tab as appropriate for the agent.

4. Click the **Content** or **Databases** tab of the Subclient Properties dialog box and Configure Subclient Content as appropriate for your agent.

5. For all agents (except QR), click the Storage Device (Data Storage Policy) tab of the Subclient Properties dialog box, then select a data storage policy to associate with this subclient from the storage policy list.

- For the DB2 and DB2 DPF *iDataAgents*, you can also change the number of data backup streams. For the DB2 DPF *iDataAgent*, the default stream threshold should be equal to the total number of database partitions for the subclient.
- For SQL Server *iDataAgents*, you can also click the Storage Device (Log Storage Policy) tab of the Subclient Properties dialog box, then select a log storage policy to associate with this subclient from the storage policy list and select the number of backup streams for transaction log backup jobs.
- For 1-Touch for Unix, it is strongly recommended that the storage policy that you select for the subclient configured for 1-Touch use a MediaAgent on a different computer. If you do this, and if the system crashes, the media will not have to be exported to another MediaAgent in order to recover the system.

6. For Oracle and DB2 *iDataAgents*, click the Backup Arguments (Oracle) or Backup Arguments (DB2, DB2 DPF) tab of the Subclient Properties dialog box and Configure Backup Arguments as appropriate for your agent. Note that the backup arguments for Informix are located on the Content tab.

7. For Migration Archiver Agents, click the **Archiving Rules** or **Rules** tab of the Subclient Properties dialog box and configure archiving rules as appropriate for your agent. In order to perform rules-based migration archiving operations, the **Disable All Rules** checkbox must be cleared.

If the File Archiver for Windows supports Data Classification, several filter-like configuration fields are defined as archiving rules and are available from the Subclient Properties (Rules) tab. If you want to define content and archiving rules based on file attributes other than volumes, size, and modified time (i.e., if you want to customize your rules), click the Advanced tab and configure as appropriate. Also, stub management options can be configured from the Stub Rule tab. See Configure Archiving Rules - File Archiver Agents for step-by-step instructions.

8. For ProxyHost and Image Level ProxyHost *iDataAgents*, click the Pre/Post Process tab of the Subclient Properties dialog box. In the **PreScan** field, type the path to the batch file/shell script file that contains those commands that are to run before each backup of the subclient, or click **Browse** to locate and select this file. For ProxyHost and Image Level ProxyHost, the file must reside on the backup host or primary host.

9. Optionally (if supported for your agent) you can:

- Add a Data Protection or Discovery Filter for a Subclient on the Filters tab.
- Configure a Subclient for Pre/Post Processing of Data Protection/Archive Operations on the Pre/Post Process tab.
- Enable Software Compression for a Subclient on the Software Compression tab of the **Storage Device** tab.
- Configure the Subclient for Data Encryption on the Encryption tab.
- Enable or Disable Operations for this subclient on the Activity Control tab.
- Configure Mailbox Stores for Auto-Discovery on the Auto-discovery tab.
- Configure the Subclient for 1-Touch on the 1-Touch Recovery tab.
- View or change the user group security associations for this subclient from the Security tab.
- Determine location from where archive logs will be backed up or deleted from the Log Destinations tab.

10. Click **OK** to save the subclient configuration. For QR Agents, this procedure is now complete. For all other agents, continue on to the next step.
11. The Backup Schedule dialog box advises you to schedule data protection operations for your new subclient. It is recommended you elect to set a schedule now. You can also associate this subclient with an All Agent Types schedule policy (which is automatically created by the system, or can be a user defined Data Protection schedule policy). If you have already associated a schedule policy at a previous level (Backup Set/Instance, Agent, Client, or Client Computer Group) the schedules defined in the Schedule Policy will be automatically applied to the new subclient. See Schedule Policy for more information.
  - o If you want to associate this subclient with an All Agent Types schedule policy, click **Associate with Generic Schedule Policy**, and then select that schedule policy from the drop-down list box. Click **OK**.
  - o If you want to associate this subclient with a specific schedule policy, click **Associate to schedule policy**, and then select the schedule policy from the drop-down list box. Click **OK**.
  - o If you have selected to define a schedule for this subclient:
    - Click **Schedule**.
    - From the Backup/Archive Options dialog box, select the type of data protection operation that you want to schedule.
    - If you want to set Advanced Backup/Archive Options, click **Advanced**.
    - After selecting the data protection type and any advanced options, click **OK**. The **Schedule Details** dialog box appears.
    - From the Schedule Details tab, select the scheduling options that you want to apply, then click **OK**.
  - o If you don't want to create a data protection schedule at this time, click **Do Not Schedule**, and then click **OK**.

This task is now complete.

## MODIFY QR POLICY PROPERTIES

It is recommended that you do not change any of the QR policy properties while the QR policy being used by an operation (e.g., QR Volume Creation, QR Volume Recovery, etc.)

*Required Capability:* Capabilities and Permitted Actions

▶ To change the properties of a QR policy:

1. Select **QR Policies** from the CommCell Browser, click **Properties** and right-click the QR policy that you wish to modify.
2. Change any of the following properties:
  - o QR policy name (Up to 32 characters).
  - o Snapshot Engine Type from the dropdown menu.
  - o Retention Policy by checking Infinite (default) or enter the number of hours/days/weeks.
  - o **Enable QR Volume Creation** check box to activate the **Copy Manager** and **Scratch Pool** menus.
  - o Copy Manager from the list of available Copy Managers that currently exist for this QR Volume.
  - o Scratch Pool from the list of available scratch pools that currently exist for this QR Volume.
  - o If this QR Policy will be used to create Recovery Points, select **Enable Recovery Points**, and specify the maximum number of Recovery Points that will be retained, up to a maximum of 32.
    - When you select **Enable Recovery Points**, you are prompted to enter a location for the cache partition for the snapshots of the QR Volume. This location cannot be on the QR Volume.
3. Once you have made your changes, click **OK**.

## SCHEDULE QR VOLUME CREATION

You can schedule QR jobs to occur with the following procedure. You will be prompted to create a schedule for the job after selecting your QR Volume creation options. See Scheduling for more information.

*Required Capability:* See Capabilities and Permitted Actions

▶ To create a QR Volume of a selected subclient:

1. From the CommCell Browser, right-click the subclient that you want to copy and then click **Create QR Volume(s)**.
2. From the QR Volume Creation window, select when you want to run the job by clicking **Run Immediately**.
3. Select **Create QR Volume(s)**.
4. Optionally, click **Advanced** to access the QR Volume Creation Advanced Options dialog box.
5. After selecting the QR Volume type and any advanced options, click **OK**.

6. From the Schedule Details (Schedule Details) dialog box, create a schedule for this operation. See Creating a Job Schedule for more information. Click **OK** to continue.
  7. Your QR Volume Creation will execute according to the specified schedule.
- 

[Back to Top](#)

# Creating a QR Volume

Topics | How To | Support | Related Topics

---

This feature/product/platform is on Extended Support in this release. See [Deprecated Features, Products, and Platforms](#) for more information.

## Overview

- Copy Only Used Blocks
- Destination Volumes
- Application Volumes

## Full QR Volume Creation

## QR Volume with Incremental Updates

- Software Compression

## Snapshot

## QR Volume Creation Job Phases

- PreSnap Phase, Snapshot Phase, and PostSnap Phase
- PreCopy Phase, QR Volume Creation Phase and PostCopy Phase
- Delete Snapshot Phase

## QR Volume Creation Considerations

- General
- QR Volume on UNIX
- QR Volume on Windows
- QR with SnapVault
- Creating a QR Volume for use with Exchange
- Creating a QR Volume for use with Oracle Database Instances
- Creating a QR Volume for use with SQL Servers
- Moving Application Data and Resources

## Related Reports

---

## OVERVIEW

QR Volumes may be created immediately from the CommCell Console, or the QR Volume creation job may be scheduled to run at a later time. It is also possible to define a schedule that creates new QR Volumes on a recurring basis, in which case the Quick Recovery Agent will automatically select the destination volumes from the Scratch Volume Pool.

QR Volume creation is not managed on a per-volume basis. Instead, you must create subclients. The volumes contained in the subclient definition will be synchronized and copied at the same point in time, to ensure the consistency of any application data sets that may span multiple volumes (for example, an Exchange 2000 Storage Group.)

---

## COPY ONLY USED BLOCKS

When using the LAN copy manager, only the used blocks of the source disk will be copied to the QR Volumes. For the QR Agent on Unix, this functionality can be turned off, so that even the empty blocks are backed up, by configuring the `dDisableDataBlock` registry key.

---

## DESTINATION VOLUMES

Destination volumes are analogous to the backup media of other agents. For any QR Volume creation job, a destination volume (from the Scratch Volume Pool) may be optionally selected for each of the primary volumes in the subclient. The destination should be at least slightly larger than the source volume. If you do not select a destination volume, the Quick Recovery Agent will do so automatically. Likewise, a destination mount point may be specified for each QR Volume. If mount points are not specified, the newly created QR Volumes will remain unmounted until you request otherwise.

## SNAPVAULT DESTINATION VOLUMES

For SnapVault, if a destination volume is not selected, the system will choose one from the scratch volume pool according to the following criteria:



- The system checks that the total size of potential destination is larger than the source.
- The system checks that the destination does not already have a replica of that source volume.
- The system checks that the remaining space on the destination volume is greater than the total size of the source.

If more than one destination meets these criteria, then the smallest destination is chosen. This does not ensure that the destination volume will not run out of space. If a destination has multiple replicas, the space used by each can grow over time. Note that you can change the default behavior for choosing destination volumes for SnapVault by changing the `sFindTargetVolUsingFreeSpace` value in the registry.

Normally, each volume in a subclient is automatically assigned to its own destination volume; however for SnapVault, there is not a one to one relationship between source and destination volumes. If you create a SnapVault subclient that contains multiple volumes, and do not specify destination volumes, the snapshots for all of the volumes could reside on one destination volume (provided there is sufficient space).

---

### APPLICATION VOLUMES

When creating QR Volumes that take advantage of the agent's application intelligence, subclients must be configured correctly to successfully recover an application volume. See the following for more information on setting up a subclient for QR Volume Creation:

- Creating a QR Volume for use with Exchange
- Creating a QR Volume for use with Oracle Database Instances
- Creating a QR Volume for use with SQL Servers

---

### FULL QR VOLUME CREATION

A full QR Volume creation operation creates a QR Volume that contains all the data associated with a subclient. QR Volumes for any subclient start with a full QR Volume Creation (i.e., all of the data is copied to the QR Volume). If you select Incremental Update as your QR Volume creation option, the first operation will be a full QR Volume creation, which becomes the baseline to which subsequent Incremental Updates are applied; in this scenario, once the Full QR Volume creation has completed, the next operation will be an incremental update.

---

### QR VOLUME WITH INCREMENTAL UPDATES

The purpose and process of creating a QR Volume with incremental updates are very similar to those of an ordinary QR Volume. However, certain restrictions apply to incremental update jobs:

- The incremental update option is not available for immediate jobs.
- Incremental update operations must be scheduled. If the associated QR policy is changed the original policy will be used until a new schedule is defined.
- Destination volumes must be selected for incremental update jobs.
- Incremental updates cannot be initiated on a previously established QR Volume.
- Only one QR Volume with incremental updates may be created for a given primary volume.
- The QR Agent must be installed on both the source client and the destination host when the source and destination volumes are on separate computers.
- The MediaAgent must be installed on the host computer of the destination volume(s).

---

### SOFTWARE COMPRESSION

Software compression is available for QR Volumes created over a LAN copy manager and is enabled or disabled at the subclient level. The data is compressed on the source client before it is copied to the destination volume. Enabling compression typically reduces the amount of network bandwidth required to copy the data; however, it introduces a computational load on the source client. Before the data is written to the destination volume, it is decompressed.

---

### SNAPSHOT

In some cases it is possible to create only a snapshot, without creating a QR Volume. See Support Information - QR Volume Creation Options for information on supported configurations. Snapshots are not backups of the data, and have unique functionality and requirements. See Snapshots for more information on the behavior of snapshots.

---

### QR VOLUME CREATION JOB PHASES

The process of creating a new QR Volume consists of seven phases:

1. PreSnap phase
2. Snapshot phase

3. PostSnap phase
4. PreCopy phase
5. QR Volume Creation phase
6. PostCopy phase
7. Delete Snapshot phase

Note that when using snapshot engines other than QSnap for Windows, the phases for QR Volume creation may vary. The phases for creating a QR Volume with QSnap for Windows are described below.

---

### **PRESNAP PHASE, SNAPSHOT PHASE, AND POSTSNAP PHASE**

In the Snapshot phase, the agent synchronizes with the applications and the operating system to ensure that all data is flushed to disk. Furthermore, it ensures that the source or primary disk image is not modified for the duration of the volume copy phase, so that the destination disk (i.e. the QR Volume) contains a consistent image.

The QR Agent will check the Scratch Volume Pool for eligible destination volumes before creating the snapshot. Ensuring that there is an eligible destination volume prevents unnecessary snapshots.

To dismount or lock the primary volume throughout the copy would require an unacceptable period of downtime for the application. Instead, the QR Agent invokes an external snapshot mechanism to create a frozen, point-in-time image of the primary volume. Once the snapshot image has been created, the application may be resumed immediately and continue to update the primary volume while the copy operation is in progress. The data is copied out of a snapshot cache. The QR Agent no longer needs to freeze the snapshot cache during the QR Volume Creation phase (see below). This reduces job failures caused by the QR Agent waiting to freeze the cache.

The use of a volume snapshot ensures that file system and application metadata remain unchanged, and therefore consistent, during the copy operation. Before creating the snapshot, the agent ensures that any supported applications have been ordered to flush buffered data and to suspend I/O to the disk (if necessary). Once the snapshot has been created, the agent signals to the application that it may resume operations. The Quick Recovery Agent is designed to accommodate a variety of external snapshot engines, including both hardware and software implementations.

The Quick Recovery Agent provides for optional PreSnap and PostSnap phases. In these phases, the agent can run a user-supplied command line or script to perform any additional processing that is required - for example, to synchronize with an application that is not yet directly supported by the agent.

---

### **PRECOPY PHASE, QR VOLUME CREATION PHASE, AND POSTCOPY PHASE**

In the QR Volume Creation phase, the Quick Recovery Agent will perform a block level copy of the data from the source disk to the destination disk (which will become the QR Volume). Depending on the QR Policy that you have selected, the copy phases will use a LAN Copy Manager to effect the data transfer. On the Windows platform, the QR Agent now zeroes out the boot sector of the destination volume so when the computer is brought back on-line, no changes will be made to the volume (keeping it consistent with the QR Volume).

The Quick Recovery Agent allows for optional PreCopy and PostCopy phases. In these phases, you may supply custom command lines or scripts, which the agent will invoke before and after the QR Volume Creation phase.

---

### **DELETE SNAPSHOT PHASE**

In the final phase, the Quick Recovery Agent deletes the snapshot that was created during the Snap phase, so that the resources are available for future QR Volume operations.

---

## **QR VOLUME CREATION CONSIDERATIONS**

---

### **GENERAL**

Before performing any QR Volume procedures for this agent, review the following information:

- When using the LAN copy manager, only the used blocks of the source disk will be copied to the QR Volumes.
- After QR Volume Creation, the target volume will become the same size as the source volume. For example, if the source volume is 5GB and the destination volume is 10GB, after a completed QR Volume Creation operation the destination volume will appear to be 5GB. Note that the extra space on NTFS volumes can be recovered by using the `expandfs.exe` tool for Windows, located on the *Resource Pack*.
- Before performing the first full backup of a volume, defragment it. This will ensure that the minimum number of extents will need to be backed up, resulting in better performance and lower storage requirements.
- For Windows 2003 R2 computers, ensure that the volume protected by QR does not have Distributed File System Replication (DFSR) configured as this may cause QR jobs to fail.
- If you encounter any locking issues on the destination volume, use the `nForceDismountToLockIfRequired` registry key to resolve the issue and unlock any locked volumes.

If you create a QR Volume with a Recovery Point and mount it to a drive letter, and then try to run `expandfs.exe` on the QR Volume, you may get a General Protection Fault error.

---

## QR VOLUME ON UNIX

- If the source volume is a slice 0 volume, the destination volume must also be a slice 0 volume. To specify the slice 0 volume to be used as the destination volume, manually select the destination volume in the QR Volume Creation Advanced Options dialog box.
- For Unix, if the slice 0 partition is empty and the disk is allocated space from slice 1 onwards, then slice 1 will be considered as slice 0 by the agent. So, from whatever slice partition you start to allocate space on the disk, that particular slice will be considered as slice 0.
- For Solaris, under conditions of heavy I/O, such as is expected during block-level backups and QR volume creation, we recommend that you do not enable the UFS logging option on the client machine.

---

## QR VOLUME ON WINDOWS

Restartability is supported for QR Volume creation on Windows platforms. See QR Volume Creation Restartability for more information.

---

## QR WITH SNAPVAULT

- Restartability is supported for QR Volume creation and update jobs for ONTAP. Restartability is not supported for QR Volume creation and update jobs on Open Systems.
- When creating a snapshot of an Oracle volume on ONTAP, you must immediately perform an incremental after the first full snapshot in order to ensure the database is in a consistent state. After this incremental, updates can be performed normally, as desired.
- If an ONTAP full backup is suspended, data added during the suspension is not automatically picked up when the job resumes. You must run an incremental backup to pick up the changes. (Note that this also applies to the first incremental backup, which is actually a full backup.)
- If you have a multi-volume subclient and the QR Volume Creation job is killed after some volumes have completed, the replicas on the destination volume will not be in synch. (The snapshots of the destination volume will be in synch.) If the transfer was incremental, the replicas will be in synch after the next successful incremental. If full, the replicas will be in place on the destination file server(s), but will not be in the database.

If a replica for the same source volume is attempted to the same destination volume having one of these replicas, the transfer will fail because the replica is already there. If this is the case, you must manually delete the SnapVault relationship or create the registry string value `sUseExistingReplica` in the `ReplicationIDAAgent` key and set it to `y`, causing the software to consider the error that the replica is already there as a successful full transfer. The replica is not updated in this case so it is not in sync with any other replicas; therefore, an incremental transfer should be run after the job completes. This registry key should be removed after the job completes to insure proper error checking.

---

## CREATING A QR VOLUME FOR USE WITH EXCHANGE

When Exchange is added to the contents of a subclient, all volumes associated with the application are added to the subclient. These Exchange volumes can be recovered at the storage group level. In order to recover at the storage group level, consider the following:

- Resource files (`.edb` or `.log` files, for example) for multiple storage groups should not exist on the same volume. This ensures that only the correct resource files are recovered when the QR Agent mounts the QR Volume. If mismatching resource files are recovered, the storage group could become corrupted.
- When you add the application to the subclient, all volumes that are associated with the Exchange data will be displayed and can be added to the subclient content.
- Generally, a full backup is recommended before creating a QR volume. See Recovering QR Volumes.
- Exchange services must be running for QR Volume creation to succeed.

---

## CREATING A QR VOLUME FOR USE WITH ORACLE DATABASE INSTANCES

It is recommended that each database instance volume be associated with its own subclient. When you add the application to the subclient, all volumes will be added to the subclient content. The volumes for the other instances should be deleted from subclient contents. See the help for instructions on modifying subclient content. This configuration has several advantages:

- Only one instance will be quiesced when performing QR operations.
- Each instance can have its own QR policy.
- Incremental updates will be faster.

If you have added or removed any data files or log files to an Oracle database, the changes need to be detected by re-adding the application to subclient contents or by creating a new subclient for the instance.

Configurations with data files, archive logs and control files from multiple instances on the same volume are not supported by the QR Agent.

Password and initialization (`initsid.ora`, `orapwsid`) files are not copied by the QR Agent to the QR Volume. If they are required for a particular recovery scenario, create these files prior to recovery.

If you plan to configure Oracle database destination volumes as unmounted, see this Note regarding recovery.

---

## CREATING A QR VOLUME FOR USE WITH SQL SERVERS

When SQL is added to the contents of a subclient, all volumes associated with the application are added to the subclient. This includes all user-defined databases and resources. System databases are filtered out of the subclient contents by the QR Agent and should not be subclient contents.

These SQL volumes can be recovered at both the instance and database level. In order to recover SQL at the database level, consider the following:

- Each database must be on its own volume(s). The volume associated with a database should not contain resources or data files for any other database.
- `.ldf` and `.mdf` files for multiple databases should not exist on the same volume. This ensures that only the desired `.ldf` and `.mdf` files are recovered when the QR Agent mounts the QR Volume. If mismatching files are recovered, the database could become corrupted.
- If you intend to configure multiple subclients for the SQL data (for example, subclients for the different database volumes), then subclients should also be created for the associated resources (`.ldf` and `.mdf` files).
- When you add the application to the subclient, all volumes that are associated with the SQL data will be displayed and can be added to the subclient content.
- If you use VSS to snap an SQL database, volume creation will succeed even when the database is corrupt or when SQL services are down. Ensure that your SQL database is not corrupt before using VSS to snap it.

---

## MOVING APPLICATION DATA AND RESOURCES

If application data is moved, the subclient contents must be redefined (i.e., the application volumes must be redetected and added to the subclient contents) and a QR Volume creation operation performed.

Once application data or resources have moved, the QR Agent cannot automatically recover the existing QR Volumes. If a QR Volume creation operation cannot be run after application data or resources have been moved, the existing QR Volumes can be mounted manually. This is not recommended; applications will not be quiesced by the QR Agent, and integrity of the recovered data cannot be guaranteed.

---

## RELATED REPORTS

---

### QR VOLUME CREATION JOB SUMMARY REPORT

The QR Volume Creation Job Summary Report provides a summary of QR volume creation operations for each client.

[Back to Top](#)

# Creating a QR Volume - How To

[Topics](#) | [How To](#) | [Support](#) | [Related Topics](#)

[Convert a QR Volume into a Primary Volume](#)

[Create a Snapshot of a Subclient](#)

[Create a QR Volume of a Subclient](#)

[Create a QR Volume that is Incrementally Updated](#)

[Delete QR Volumes](#)

[Mount and Unmount QR Volumes](#)

[Create or Remove CIFS/NFS Shares for SnapVault Snapshots](#)

[Schedule QR Volume Creation](#)

[Transfer Data from a QR Volume to a New Primary Volume](#)

---

## CONVERT A QR VOLUME INTO A PRIMARY VOLUME

Once a QR Volume has been recovered, you may decide to continue running the application from that volume on a permanent basis. In that case, you will probably want to reconfigure QR Agent to treat that volume as the primary volume.

▶ To convert a QR Volume into a primary volume:

1. Use the QR Volume Details screen to identify the QR Policy and Scratch Volume Pool associated with the QR Volume.
2. Go to the Scratch Volume Pool in the CommCell Browser and find the destination volume.
3. Right-click on the destination volume and select **Release** from the drop-down menu. This will cause the QR Volume association to be removed from the

volume, and the destination volume itself will be removed from the Scratch Volume Pool.

4. Reconfigure the QR Agent subclient content to reflect the new primary volume, and remove the old primary volume from the subclient.
- 

## CREATE A SNAPSHOT OF A SUBCLIENT

### Before you Begin

- Ensure there are no handles on the destination volumes before creating the QR Volume.
- See Create a VSS Hardware Snapshot of a Subclient if you are creating a VSS Hardware snapshot.

*Required Capability:* See Capabilities and Permitted Actions

▶ To create a snapshot of a selected subclient:

1. From the CommCell Browser, right-click the subclient that you want to snap and then click **Create QR Volume(s)**.
  2. From the QR Volume Creation window, select when you want to run the job by clicking **Run Immediately**.
  3. Select **Snapshot**.
  4. Optionally, click **Advanced** to access the QR Volume Creation Advanced Options dialog box.
  5. After selecting the QR Volume type and any advanced options, click **OK**. The QR Volume operation will begin and you can track the progress of the job in the Job Controller window.
- 

## CREATE A QR VOLUME OF A SUBCLIENT

### Before you Begin:

- Make sure you check QR Volume Creation Considerations for your agent.

Create QR Volume(s) allows you to secure data immediately without having to wait for the scheduled QR Volume Creation time. This capability can be useful if:

- You have some particularly valuable volumes that you need to secure immediately.
- You want to create a QR Volume that is not routinely secured through a QR Volume schedule.

*Required Capability:* See Capabilities and Permitted Actions

▶ To create a QR Volume of a selected subclient:

1. From the CommCell Browser, right-click the subclient that you want to copy and then click **Create QR Volume(s)**.
  2. From the QR Volume Creation window, select when you want to run the job by clicking **Run Immediately**.
  3. Select **Create QR Volume(s)**.
  4. Optionally, click **Advanced** to access the QR Volume Creation Advanced Options dialog box.
  5. After selecting the QR Volume type and any advanced options, click **OK**. The QR Volume operation will begin and you can track the progress of the job in the Job Controller window.
- 

## CREATE A QR VOLUME THAT IS INCREMENTALLY UPDATED

*Required Capability:* See Capabilities and Permitted Actions

▶ To create a QR Volume that is incrementally updated:

1. From the CommCell Browser, right-click the subclient that you want to copy and then click **Create QR Volume(s)**.
2. From the QR Volume Creation window, select when you want to run the job by clicking **Schedule**. It is important to note that Incrementally update QR Volume(s) is only available when Schedule is selected.
3. Select the type of operation that you want to initiate. Select **Incrementally update QR Volume(s)**.
4. Optionally, click **Advanced** to access the QR Volume Creation Advanced Options dialog box.

One of the available Advanced Options when you create a QR Volume that is incrementally updated (and when the associated QR Policy has Recovery Points enabled) is the option to mount the latest Recovery Point to a specified mount point. This option, Mount Destination Snap, allows access to the destination snap during the next incremental QR Volume creation job when the destination volume is normally not accessible. This provides uninterrupted access to the destination even during the copy phase of the next incremental QR Volume creation job.

5. After selecting the QR Volume type and any advanced options, click **OK**.

6. From the Schedule Details (Schedule Details) dialog box, create a schedule for this operation. See [Creating a Job Schedule](#) for more information. Click **OK** to continue.
  7. Your QR Volume Creation will execute according to the specified schedule. The first operation will be a full QR Volume creation. Subsequent updates will be made incrementally according to the schedule you have created.
- 

## DELETE QR VOLUMES

When a QR Volume is no longer needed, you can delete it to release the volume for other uses.

▶ To delete a QR Volume:

1. Go to the QR Volume Browse window.
2. Select the volume.
3. Click **Delete**.

The following restrictions apply to QR Volume deletion:

- A QR Volume cannot be deleted when the In Use flag is set. This flag is set automatically by a QR Volume recovery operation to prevent the QR Volume from being pruned or deleted while an application is using it. You can clear the flag on the QR Volume Details screen.
  - A QR Volume cannot be deleted while it is mounted on a host. Remove all mount points from the QR Volume before you delete it.
  - In addition to deleting QR Volumes, they can be released. Releasing a QR Volume allows it to be selected for future operations (overwriting any data existing on the volume).
- 

## MOUNT AND UNMOUNT QR VOLUMES

QR Volumes may be mounted or unmounted like any other disk volume. The QR Volume can be mounted on a different host, provided that the host has the QR Agent or MediaAgent software installed.

You can add or remove mount points on a QR Volume from the QR browser or from Volume Explorer.

▶ To add or remove mount points from the QR browser:

1. Open/go to the QR Volume Browse window.
2. Select a QR Volume.
3. Right-click **Details**, and go to the Mount Points tab.
4. Select the host machine where the volume will be mounted.

▶ To add or remove mount points from Volume Explorer:

1. From Volume Explorer, select the QR Volume.
2. Use the Mount Volume and Unmount Volume controls.

See [Browse Available QR Volumes](#) for more information on [Browsing QR Volumes](#).

- SnapVault snapshots cannot be mounted but you can create a CIFS (Windows) or NFS (Unix) share to the QR source or destination snapshot. See [Create CIFS/NFS Shares for SnapVault Snapshots](#) for more information.
- 

## CREATE OR REMOVE CIFS/NFS SHARES FOR SNAPVAULT SNAPSHOTS

For SnapVault snapshots CIFS (Windows) or NFS (Unix) shares to the QR source or destination snapshot can be created.

You can create or remove CIFS/NFS shares to the QR source or destination snapshots from the QR browser.

▶ To create or remove CIFS/NFS shares from the QR browser:

1. Open/go to the QR Volume Browse window.
2. Select a SnapVault snapshot.
3. Right-click **Details**, and go to the CIFS/NFS tab.
4. Select the host machine where the share will be created.
5. Click **ADD**, and specify the share name in the Add Share name dialog box.

- For windows, you can specify any share name but in case of unix, the export name is always set to the full path

on the file server and cannot be modified.

- See [View a List of Available QR Volumes](#) for more information on [Browsing QR Volumes](#).
- 

## SCHEDULE QR VOLUME CREATION

You can schedule QR jobs to occur with the following procedure. You will be prompted to create a schedule for the job after selecting your QR Volume creation options. See [Scheduling](#) for more information.

*Required Capability:* See [Capabilities and Permitted Actions](#)

▶ To create a QR Volume of a selected subclient:

1. From the CommCell Browser, right-click the subclient that you want to copy and then click **Create QR Volume(s)**.
  2. From the QR Volume Creation window, select when you want to run the job by clicking **Run Immediately**.
  3. Select **Create QR Volume(s)**.
  4. Optionally, click **Advanced** to access the QR Volume Creation Advanced Options dialog box.
  5. After selecting the QR Volume type and any advanced options, click **OK**.
  6. From the Schedule Details (Schedule Details) dialog box, create a schedule for this operation. See [Creating a Job Schedule](#) for more information. Click **OK** to continue.
  7. Your QR Volume Creation will execute according to the specified schedule.
- 

## TRANSFER DATA FROM A QR VOLUME TO A NEW PRIMARY VOLUME

In certain disaster recovery scenarios, application data can be accessed from the associated QR Volume. For example, this may be necessary if the storage subsystem containing the primary volumes becomes unavailable for any reason.

Once the primary storage subsystem is available, in most cases it is desirable to transfer the contents of the temporary QR Volume to a new primary volume on the primary storage subsystem. This is desirable because the storage subsystems containing your primary volumes are usually significantly more powerful and/or more expensive than the storage used for temporary QR Volumes.

▶ To transfer your data from a temporary QR Volume to a New Primary Volume:

1. Convert your QR Volume(s) to primary volumes using the procedure outlined in [Converting a QR Volume into a Primary Volume](#).
2. Update your subclient content definitions.
3. Partition the (new) primary storage device into volumes and add them to the appropriate Scratch Volume Pools. It may be necessary to re-detect the devices before they can be added.
4. Run one or more Create QR Volume jobs to copy data from the temporary volumes to the permanent volumes.
5. Convert the new QR Volume(s) (located on the new permanent storage) into primary volumes, again using the procedure outlined above.

This two-stage procedure may be especially valuable in LAN-only environments, where it might not be otherwise possible to mount the QR Volumes directly on the application server.

---

[Back To Top](#)

# Recovering QR Volumes

Topics | How To | Related Topics | Use Cases: QR Disaster Recovery

---

This feature/product/platform is on Extended Support in this release. See [Deprecated Features, Products, and Platforms](#) for more information.

Overview

Quick Recovery

- [Recovering Application Data](#)
- [Recovering Individual Volumes](#)

[Recovering Consistent Recovery Points](#)

Copyback

- [Snapshot Copyback Considerations](#)

Recovery Destinations

- [In-Place Recovery](#)
- [Out-of-Place Recovery](#)
- [Cross-Platform Recovery](#)

[Recovery Considerations for This Agent](#)

[Related Reports](#)

---

## OVERVIEW

The Quick Recovery Agent supports the following types of recovery:

- [Quick Recovery](#)
- [Copyback](#)

The ultimate purpose of QR Volumes is to provide for rapid recovery of critical data. The QR Agent provides automated tools to assist you in the recovery process.

For the Quick Recovery Agent, recovery operations can be performed from the client and agent levels in the CommCell Browser.

---

## SNAPVAULT RESTORES

SnapVault does not support Quick Recovery or Copyback. Data is browsed and restored at the file or volume level. See [Restore Data - SnapVault](#) for an overview of restoring data using SnapVault.

---

## QUICK RECOVERY

During a quick recovery operation, the source volume is unmounted, and then the QR Volume is automatically mounted in its place (with the same mount point). The mount point for the source volume must exist in order to swap the mount points during a quick recovery. If the source volume is corrupted or not available, mount another volume to the source volume mount point before performing the quick recovery. Specifically, when a recovery operation is executed:

- any applications that use the volume (e.g., SQL, Exchange) are suspended.
- the primary (source) volume is dismounted.
- the QR Volume is mounted in its place.
- the applications are resumed.

In addition to Quick Recovery, the QR Agent also provides a Copyback feature.

Consistent Recovery Points, created through the Quick Recovery Agent and QSnap, can be recovered using the Copyback feature. See [Recover Consistent Recovery Points](#).

---

## RECOVERING APPLICATION DATA

When restoring application data, steps must be taken to suspend and resume the application gracefully while ensuring the integrity of the application data



structures.

See Overview - Quick Recovery Agent for a list of applications that the QR Agent can automatically recover.

To recover a supported application volume, the respective *iDataAgents* must be installed on the client together with the QR Agent. Browse the backup data on the appropriate *iDataAgent*, and select the snapshot version of the application data object (e.g., Exchange 2000 Storage Group) that you wish to recover. The QR Agent will automatically suspend the application, dismount the old primary volume(s), and remount the appropriate QR Volume(s) in their place.

The QR Volumes will be recovered in the following states, depending upon the type of volume recovered:

- **Exchange Server**

The stores are recovered in an unmounted state. Exchange services must be running for Quick Recovery of Exchange data to succeed.

- **SQL Server**

With QSnap, when SQL data is being recovered, the user will be asked if the database should be put in the loading state. If the user selects **Yes**, the database will be in the loading state. If the user selects **No**, the database will be online.

With VSS, SQL data recovery will always leave the database in the loading state.

- **Oracle**

The databases are mounted, but not open.

These states allow for maximum flexibility once the QR Agent has recovered the volume(s). For example, because the volumes are not mounted to be active, logs can still be applied to update the volumes.

If a QR volume is created before a full backup, then the gap of transactions between the two could cause an out-of-sequence log restore. Because of this, log restores are typically not allowed if a full backup was performed after a QR Volume was created.

---

## RECOVERING INDIVIDUAL VOLUMES

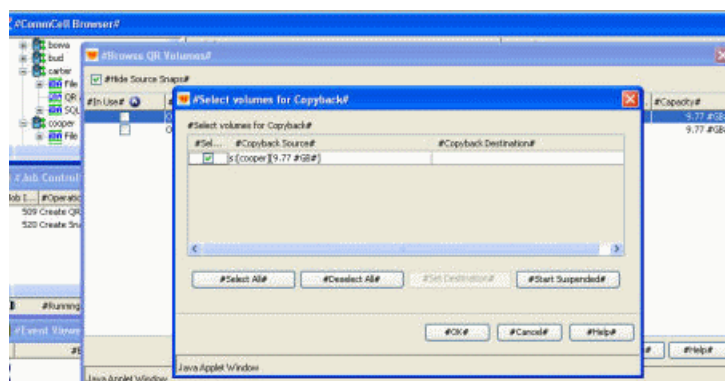
To recover an individual volume, use the Browse QR Volumes window. As with the application-specific recovery process, the primary volume will be dismounted and the QR Volume will be mounted in its place. No special steps are taken to synchronize applications in this case; synchronization must be done manually.

### NOTES

- To recover Oracle database destination volumes that you had configured as unmounted, you must run the `fsck` command first.

## RECOVERING CONSISTENT RECOVERY POINTS

Consistent Recovery Points created with the QR Agent are recovered in the Browse QR Volumes window through the Copyback feature; follow the Copy Back a QR Volume procedure. Using Copyback, you can recover both mounted and unmounted Recovery Points.



For more information about this feature, see Recovery on the Recovery Points page, and Copyback and Snapshot Copyback Considerations on this page.

## COPYBACK

A copyback operation copies a QR Volume or snapshot to another volume. The destination volume is selected when the copyback operation is started in the QR Volume Browse window. Multiple QR Volumes and/or snapshots can be copied back by selecting multiple volumes in the QR Volume Browse window. The data is copied directly from the QR Volume; no snapshot is necessary for this operation. When copying volumes using copyback, note the following:

- Copyback is supported only with LAN copy manager configurations.
- Software compression for the copyback operation can be enabled with the `QRCompressionSupported` registry key.
- On Windows, you can copy back to the original source volume regardless of its size. (This is not true for Unix.)

---

## SNAPSHOT COPYBACK CONSIDERATIONS

- Copyback is not supported for Volume Shadow Copy Services (VSS) software snapshots or for VSS snapshots created with a hardware provider, but copyback is supported for snapshots of the source created by QSnap.
  - EMC SnapView snapshots that you want to copy back must be activated. If the snapshot is not activated, the copyback will fail. You can activate the snapshot from the QR Volume Browse window.
  - For QR Volumes created using a File System Snapshot in Solaris, you must first delete the source snapshot before you can copy back the QR Volume to the source; if you do not, the source volume will not be able to be unmounted for the copyback operation. To delete the source snapshot, unmount and run the `fssnap -d` command.
  - Do **not** select the snapshot cache volume as the copyback destination even though it will show up as a possible copyback destination if it is larger than the source volume. During copyback operations, the destination volume needs to be unmounted, but a snapshot cache volume must remain mounted during the entire copyback operation. Selecting the cache volume as the copyback destination will cause the job to fail.
- 

## RECOVERY DESTINATIONS

By default, the Quick Recovery Agent recovers volumes to the client computer from which it originated. This is comparable to an in-place restore on other agents. The Quick Recovery Agent is a volume-level agent and has unique recovery destinations described below.

You can also manually mount the QR Volume or copy back the volume to another client computer in the CommCell. Keep in mind the following considerations when mounting or copying back volumes to different computers:

- Manually mounting or copying back a volume is not intended to be used for recovering volumes into a live production environment because applications will not be quiesced. Volumes will not be ready for use by their associated applications.
- Manually mounting and/or Copying back a QR Volume only provides an alternative means to access data on your QR Volumes.

The following section enumerates the types of recovery destinations that are supported by the Quick Recovery Agent. See Restore/Recover/Retrieve Destinations - Support for a list of Agents supporting each restore destination type.

---

## IN-PLACE RECOVERY

Quick Recovery operations are always in-place. The QR Volume is assigned the mount point of the source volume (after automatically unmounting the source volume). Therefore, Quick Recovery always recovers to the same operating system, client computer, etc.

---

## OUT-OF-PLACE RECOVERY

Quick Recovery is not supported out-of-place. Copyback and manual mounting of QR Volumes provide the basic means for accessing data out-of-place, taking into account the considerations listed above.

---

## CROSS-PLATFORM RECOVERY

Because Quick Recovery operations are always to the same computer, cross-platform recovery is not supported.

When mounting QR Volumes or copying back QR Volumes to a different platform consider the following:

- Copyback between Unix and Windows is not supported.
  - The computer must support the file system of the volume to be mounted or copied back.
  - On the Windows Platform, when copying back or mounting a volume to a client using a different version of Windows, see Cross-Platform Restores for a list of considerations.
- 

## RECOVERY CONSIDERATIONS FOR THIS AGENT

To avoid common problems, review the following before starting a recovery operation:

- The last QR Job must have completed successfully in order to perform a successful recovery. If the last incremental update or QR Volume creation operation failed, the data will be physically recovered but it may be invalid.
  - It is recommended that you delete or disable schedules associated with the subclient that is to be recovered.
  - Ensure that there are no handles present on the source and destination volumes.
- 

## RELATED REPORTS

---

### QR RECOVERY JOB SUMMARY REPORT

The QR Recovery Job Summary Report provides a summary of QR recovery operations for each client.

---

[Back to Top](#)

# Recovering QR Volumes - How To

[Topics](#) | [How To](#) | [Related Topics](#) | [Use Cases: QR Disaster Recovery](#)

---

[Apply Transaction Logs to an Oracle Database on a Recovered QR Volume with Archive Log Backups](#)

[Apply Transaction Logs to an Oracle Database on a Recovered QR Volume without Archive Log Backups](#)

[Convert a QR Volume into a Primary Volume](#)

[Copy Back a QR Volume](#)

[Recover a QR Application Volume](#)

[Recover a QR Volume](#)

[Transfer Data from a QR Volume to a New Primary Volume](#)

---

## APPLY TRANSACTION LOGS TO AN ORACLE DATABASE ON A RECOVERED QR VOLUME WITH ARCHIVE LOG BACKUPS

Before starting restore operations, review these prerequisites and operational procedures.

- Review the information in [Using Transaction/Archive Logs with Recovered QR Volumes](#).
- You have installed the Oracle iDataAgent and QR Agent on the same client.
- You have a snapshot engine, such as QSnap.
- You have performed any one of the following tasks for the Oracle instance:
  - Full QR Volume Creation only
  - QR Incremental Updates using a schedule
  - Full QR Volume Creation with logs backup
  - QR Incremental Updates as well as logs backup

*Required Capability:* See [Capabilities and Permitted Actions](#)

▶ To recover a database from a QR Volume with archive log backups.

1. Perform a QR Recovery to recover the point-in-time Oracle data from your QR Volumes. Refer to [Recovering QR Volumes](#) for procedures on performing a QR Recovery.
2. After recovering the point-in-time Oracle data from your QR Volumes, access the CommCell Console, right-click the instance of the Oracle iDataAgent, click **All Tasks** and then click **Restore**.
3. From the Oracle Restore Options (General) dialog box, select **Restore Archive Log** and clear the **Restore Data** checkbox. Click **Advanced**.
4. From the Restore tab of the Oracle Advanced Restore Options dialog box, under the Restore Archive Log - By Log Time options, clear the **Start** checkbox then click **End** and select the date and time that the most recent archive log backup completed. Click **OK**.
5. When restoring encrypted data, refer to [Data Encryption](#).
6. Start or schedule the restore. The restore job will begin and you can track the status of the job in the Job Controller window. After the job completes, all the archive logs are restored to the specified point-in-time.
7. After the archive logs have been restored, right-click the instance of Oracle iDataAgent, click **All Tasks** and then click **Restore**.
8. From the Oracle Restore Options (General) dialog box, select **Recover** and clear the **Restore Data** checkbox. Click **Advanced**.
9. From the Recover tab of the Oracle Advanced Restore Options dialog box, under the Recover options, click **Point in Time** and select the date and time equal to the last logs backup time. Click **OK**.
10. Start or schedule the restore. The recovery job will begin and you can track the status of the job in the Job Controller window. After the job completes, the database will be recovered to the specified time, and the state of the Oracle database will change from MOUNTED to OPEN. This will result in all the archived logs being applied towards the instance.
 

If the database recovery fails with an error indicating that it was unable to find an archive log then continue on with this procedure, otherwise this task is complete.
11. Right-click the instance of Oracle iDataAgent, click **All Tasks** and then click **Restore**.
12. From the Oracle Restore Options (General) dialog box, select **Recover** and clear the **Restore Data** checkbox. Click **Advanced**.

13. From the Recover tab of the Oracle Advanced Restore Options dialog box, under the Recover options, click **Point in Time** then select the date and time to which you want to recover the database.
14. From the Options tab of the Oracle Advanced Restore Options dialog box, select **No Re-do Logs** and then click **OK**.
15. Optionally for Oracle RAC, from the Oracle Restore Options (Stream) dialog box, set the instance restore order per stream allotment.
16. Start or schedule the restore. The recovery job will begin and you can track the status of the job in the Job Controller window. After the job completes, the Oracle database will be recovered to the specified time, and the state of the Oracle database will change from `MOUNTED` to `OPEN`.

## APPLY TRANSACTION LOGS TO AN ORACLE DATABASE ON A RECOVERED QR VOLUME WITHOUT ARCHIVE LOG BACKUPS

### Before You Begin

Before starting restore operations, review these prerequisites and operational procedures.

- Review the information in Using Transaction/Archive Logs with Recovered QR Volumes.
- You have installed the Oracle iDataAgent and QR Agent on the same client.
- You have performed any one of the following tasks for the Oracle instance:
  - Full QR Volume Creation only
  - QR Incremental Updates using a schedule
  - Full QR Volume Creation with logs backup
  - QR Incremental Updates as well as logs backup

*Required Capability:* See Capabilities and Permitted Actions

▶ To recover a database from a QR Volume without archive log backups.

1. From the CommCell Console, right-click the instance of Oracle iDataAgent, click **All Tasks** and then click **Browse Backup Data**.
2. Select your browse options and click **OK** to continue. See Browsing Data for a list of customized browse operations and their step-by-step instructions. If you accept all defaults, you will be browsing the latest backups for the selected data.
3. From the Browse window, expand the Snapshot of your Oracle instance, select the entire instance and then click **Recover All Selected**.
4. From the Specify Database State dialog box, select **NO**. The databases will be restored in the `OPEN` state.
5. The recovery job will begin and you can track the status of the job in the Job Controller window. After the job completes, the Oracle database will be recovered to the last successful QR Volume update, and the state of the Oracle database will change from `MOUNTED` to `OPEN`. Archive logs cannot be applied after this point.

## CONVERT A QR VOLUME INTO A PRIMARY VOLUME

Once a QR Volume has been recovered, you may decide to continue running the application from that volume on a permanent basis. In that case, you will probably want to reconfigure QR Agent to treat that volume as the primary volume.

▶ To convert a QR Volume into a primary volume:

1. Use the QR Volume Details screen to identify the QR Policy and Scratch Volume Pool associated with the QR Volume.
2. Go to the Scratch Volume Pool in the CommCell Browser and find the destination volume.
3. Right-click on the destination volume and select **Release** from the drop-down menu. This will cause the QR Volume association to be removed from the volume, and the destination volume itself will be removed from the Scratch Volume Pool.
4. Reconfigure the QR Agent subclient content to reflect the new primary volume, and remove the old primary volume from the subclient.

## COPY BACK A QR VOLUME

▶ To copy back a QR Volume:

1. From the CommCell Browser, right-click the QR Agent whose volumes (including Recovery Points and snapshots) you want to browse, click **All Tasks**, and then click **Browse** from the short-cut menu.
2. From the Browse QR Volumes window, right-click the volume (or Recovery Point or snapshot) that you want to copy back, and then click **Copyback** to open the Select Volumes for Copyback dialog box. (If you want to copy back multiple QR Volumes, highlight all of the volumes you want to copy back by control-clicking the QR Volumes.)

3. From the Select Volumes for Copyback dialog box, assign a destination volume to each QR Volume you want to copy back. The destination volume must be larger than the QR Volume.
  4. To assign a destination volume, click on the QR Volume, and then click the **Set Destination** button to open the Select Destination for Copyback dialog box.
  5. From the list of volumes in the Select Destination for Copyback dialog box, select the volume that is to be the copyback destination. (You can see a list of volumes on a different host by selecting a different computer in the dropdown Host: list.) Click the volume you want as the copyback destination and click **Select**.
  6. Repeat steps 4 and 5 for any QR Volumes that do not have a destination volume for copyback in the Select Destination for Copyback dialog box.
  7. Ensure that all volumes you want to copy back are checked in the Select column.
  8. Click **OK** to copy back the volumes. The system displays a progress bar and copies back the volumes. You can track the progress of the recover operation from the Job Controller window.
- 

## RECOVER A QR APPLICATION VOLUME

*Required Capability:* See Capabilities and Permitted Actions

▶ To browse and recover a QR Application Volume:

1. From the CommCell Browser, right-click the *iDataAgent* whose data you want to browse, click **All Tasks** and then click **Browse** from the short-cut menu.
2. From the Browse window, select the Snapshot data and/or Snapshot directories (i.e., Snapshot1) that you want to recover. Right-click the selected snapshot or snapshot directory and click **Recover**.
3. A prompt appears, warning you that the recover operation will unmount the selected QR Volume and replace it with the original source mount point. If you want to proceed, click **Yes**.

The system displays a progress bar and recovers the volume. You can track the progress of the recover operation from the Job Controller window.

4. If you are recovering a database volume, you will be asked if you want to leave the database in the Loaded state.
    - o If you want to apply logs to the database after the recovery, click **Yes**.
    - o If you want to mount the database after the recovery, click **No**. You will not be able to apply logs after the recovery if you click No.

With Oracle, selecting Yes will leave the database in the mounted state. If you want to put the Oracle database in OPEN (READ WRITE) mode after the recovery, click **No**. You will not be able to apply logs after the recovery if you click No.
- 

## RECOVER A QR VOLUME

*Required Capability:* See Capabilities and Permitted Actions

▶ To browse and recover a QR Volume:

1. From the CommCell Browser, right-click the QR Agent whose volume you want to browse, click **All Tasks** and then click **Browse** from the short-cut menu.
2. From the Browse QR Volumes dialog box, right-click the volume that you want to recover and then click **Recover**.
3. A prompt appears, warning you that the recover operation will unmount the original source volume and replace it with the selected QR Volume. If you want to proceed, click **Yes**.

The system displays a progress bar and recovers the volume. You can track the progress of the recover operation from the Job Controller window.

When the selected QR Volume is recovered, it is ready for use.

---

## TRANSFER DATA FROM A QR VOLUME TO A NEW PRIMARY VOLUME

In certain disaster recovery scenarios, application data can be accessed from the associated QR Volume. For example, this may be necessary if the storage subsystem containing the primary volumes becomes unavailable for any reason.

Once the primary storage subsystem is available, in most cases it is desirable to transfer the contents of the temporary QR Volume to a new primary volume on the primary storage subsystem. This is desirable because the storage subsystems containing your primary volumes are usually significantly more powerful and/or more expensive than the storage used for temporary QR Volumes.

▶ To transfer your data from a temporary QR Volume to a New Primary Volume:

1. Convert your QR Volume(s) to primary volumes using the procedure outlined in *Converting a QR Volume into a Primary Volume*.
2. Update your subclient content definitions.

3. Partition the (new) primary storage device into volumes and add them to the appropriate Scratch Volume Pools. It may be necessary to re-detect the devices before they can be added.
4. Run one or more Create QR Volume jobs to copy data from the temporary volumes to the permanent volumes.
5. Convert the new QR Volume(s) (located on the new permanent storage) into primary volumes, again using the procedure outlined above.

This two-stage procedure may be especially valuable in LAN-only environments, where it might not be otherwise possible to mount the QR Volumes directly on the application server.

---

[Back To Top](#)

# Restore Data - SnapVault

Topics | How To | Related Topics

---

- Overview
  - Restore Considerations for SnapVault Products
  - Restore Destinations
    - In-Place Restore
    - Out-of-Place Restore
    - Cross-Platform Restores
    - Restore to Network Drive/NFS-Mounted File System
- 

## OVERVIEW

The following page describes the agent-specific restore options. Additional restore options are accessible from the Related Topics menu.

ONTAP SnapVault for the Quick Recovery Agent and OSSV SnapVault for the Quick Recovery Agent allow you perform the following types of restore operations:

- **File Level Restore** - You can browse the files/folders in your SnapVault Snapshots and Replicas and select the files/folders you want to restore.
- **Volume Level Restore** - You can restore volumes you have backed up by selecting all contents for a volume; do this by selecting the top level node in the Browse window.

## NOTES

The system does not allow cross-OS volume level restores.

---

## RESTORE CONSIDERATIONS FOR SNAPVAULT PRODUCTS

Before performing any restore procedures for SnapVault products, review the following information:

- For OSSV SnapVault, the OSSV Client must be added to the NDMP Server list in order to perform restore operations.
  - For OSSV SnapVault on Unix, only CXBF configured volumes are available for restore in Restore Options.
  - In the Job Controller Window, the progress indicator for restore stays at 0% until the job is finished, at which time the indicator changes to 100%. This occurs whether the job is successful or failed.
  - To restore an iSCSI snapshot volume, take the volume offline before performing the restore. Otherwise the restore will fail.
  - Application level restore does not apply to SnapVault products. Snapshots/Replicas should be restored through the QR Browse window as described in the procedures below.
  - **File Level Restore**
    - When restoring data to a computer in a different domain:
      - If two domains have the correct trust relationship established, the ACLs will be restored to the new domain.
      - If the two domains do not have any trusts, the old ACLs cannot be restored to the new domain even if the accounts have the same names. This is because the SIDs for the name will be different in each domain, and, therefore, are not recognized.
    - If a specified destination path does not exist on the file server, the full path will be created in root. For example, if the root is `/vol/vol10` and you specify `vol1/S3/qtreen1/test` as the destination path and it does not exist on the destination file server, then the files are restored under `/vol/vol10/vol1/S3/qtreen1/test`.
- 

## RESTORE DESTINATIONS

By default, SnapVault products for the Quick Recovery Agent restore a volume or file(s) to the client from which it originated; this is referred to as an in-place restore. If desired, you can also restore the data to a different client or file server. Keep in mind the following considerations when performing such restores:

- The destination client must reside in the same CommCell as the client whose data was replicated.
- Note that when you perform restores other than in-place restores, the restored data assumes the rights (i.e., permissions) of the parent directory.

The following section enumerates the types of restore destinations that are supported by ONTAP SnapVault for the Quick Recovery Agent and OSSV SnapVault for the Quick Recovery Agent. See [Restore/Recover/Retrieve Destinations - Support](#) for a list of Agents supporting each restore destination type.

---

## IN-PLACE RESTORE

- Same path/destination

---

### OUT-OF-PLACE RESTORE

- Same path/destination
- Different path/destination

---

### CROSS-PLATFORM RESTORES

- Same Operating System - Different Version
- Different Operating System

Cross-Platform, File Level restores are supported as follows:

- Windows to Unix
- Unix to Windows
- ONTAP to Unix
- ONTAP to Windows

Cross-Platform, Volume Level restores are not supported.

---

### RESTORE TO NETWORK DRIVE/NFS-MOUNTED FILE SYSTEM

For File Level restores only, you can also Restore to a Network Drive/NFS-Mounted File System. This is similar to an out-of-place restore except that:

- The restored data passes through the client computer to the mapped NFS mount.
- The computer that hosts the share or mounted file system need not be another client within the CommCell.

[Back to Top](#)

---

## Restore Data - How To - SnapVault

[Topics](#) | [How To](#) | [Related Topics](#)

---

[Browse and Restore \(SnapVault\)](#)

[Restore Out-of-Place \(SnapVault\)](#)


---

### BROWSE AND RESTORE (SNAPVAULT)

#### Before You Begin:

- Review Restore Data - SnapVault.

*Required Capability:* See Capabilities and Permitted Actions

 To browse and restore data:

1. From the CommCell Browser, right-click the QR Agent, click **All Tasks** and then click the available **Browse** command.
2. From the QR Volume Browse window, right-click the SnapVault Snapshot or Replica you are restoring from and click **Browse**.
3. This opens a File level Browse window. From the File level Browse window, Select Objects From the Browse Window for Restore.

#### NOTES

To Perform a volume level restore, simply select the top node in the Browse window for the volume.

4. From the agent's Restore Options dialog box, select the restore options that you want to use. When you accept all the default settings, you will be restoring the selected data to its original location.
5. After completing your selections, you can either start an immediate restore or schedule the restore.
  - If you want to schedule the job, click the Job Initiation tab from the Restore Options dialog box, click **Schedule**, and enter your selections in the Schedule Details (Schedule Details) dialog box. Clicking **OK** from this dialog box saves your schedule. See Scheduled Data Recovery Operations for an overview of this feature.
  - If you want to run the job now, accept or click **Run Immediately** in the same tab and then click **OK**.

While the job is running, you can right-click the job in the Job Controller and select **Detail** to view information on the job.



After the data has been restored, you will see a job completion message in the Job Controller and Event Viewer.

---

## RESTORE OUT-OF-PLACE (SNAPVAULT)

*Required Capability:* See Capabilities and Permitted Actions

▶ To browse and restore data:

1. From the CommCell Browser, right-click the QR Agent, click **All Tasks** and then click the available **Browse** command.
2. From the QR Volume Browse window, right-click the SnapVault Snapshot or Replica you are restoring from and click **Browse**.
3. This opens a File level Browse window. From the File level Browse window, Select Objects From the Browse Window for Restore.

### NOTES

To Perform a volume level restore, simply select the top node in the Browse window for the volume.

4. From the agent's Restore Options dialog box, select the restore options that you want to use. Then, select the restore destination:
  - **Restore to a different location** - Select this option to restore data to a different path and or client.
  - **Destination Host** - By default, data is restored to the same host from which it was backed up. To change the destination host, select one from the list.
  - **Destination Path** - By default, data is restored to the same host from which it was backed up. Use this space to enter (or browse to) the path on the destination host to which the data will be restored. If the specified path does not exist, it will be created during the restore process.
5. After completing your selections, you can either start an immediate restore or schedule the restore.
  - To set a delayed starting time or to schedule a series of restores, click the Job Initiation tab from the Restore Options dialog box, click **Schedule**, and enter your selections in the Schedule Details (Schedule Details) dialog box. Clicking **OK** from this dialog box saves your schedule. See Scheduled Data Recovery Operations for an overview of this feature.
  - If you want to run the job now, accept or click **Run Immediately** in the same tab and then click **OK**.

While the job is running, you can right-click the job in the Job Controller and select **Detail** to view information on the job.

After the data has been restored, you will see a job completion message in the Job Controller and Event Viewer.

---

[Back to Top](#)

# Restore Data - ONTAP Snapshot

Topics | Related Topics

---

- Overview
  - Restore Considerations for ONTAP Snapshot
  - Restore Destinations
    - In-Place Restore
    - Out-of-Place Restore
    - Cross-Platform Restores
    - Restore to Network Drive/NFS-Mounted File System
- 

## OVERVIEW

The following page describes the agent-specific restore options. Additional restore options are accessible from the Related Topics menu.

ONTAP Snapshot for the Quick Recovery Agent allows you perform the following types of restore operations:

- **File Level Restore** - You can browse the files/folders in your snapshots and select the files/folders you want to restore.
- **Volume Level Restore** - You can restore volumes you have backed up by selecting all contents for a volume; do this by selecting the top level node in the Browse window.

## NOTES

The system does not allow cross-OS volume level restores.

---

## RESTORE CONSIDERATIONS FOR ONTAP SNAPSHOT

Before performing any restore procedures for ONTAP Snapshot, review the following information:

- In the Job Controller Window, the progress indicator for restore stays at 0% until the job is finished, at which time the indicator changes to 100%. This occurs whether the job is successful or failed.
  - To restore an iSCSI snapshot volume, take the volume offline before performing the restore. Otherwise the restore will fail.
  - Application level restore does not apply to ONTAP Snapshot. Snapshots should be restored through the QR Browse window as described in the procedures below.
  - **File Level Restore**
    - When restoring data to a computer in a different domain:
      - If two domains have the correct trust relationship established, the ACLs will be restored to the new domain.
      - If the two domains do not have any trusts, the old ACLs cannot be restored to the new domain even if the accounts have the same names. This is because the SIDs for the name will be different in each domain, and, therefore, are not recognized.
    - If a specified destination path does not exist on the file server, the full path will be created in root. For example, if the root is `/vol/vol10` and you specify `vol/S3/qtrees1/test` as the destination path and it does not exist on the destination file server, then the files are restored under `/vol/vol10/vol/S3/qtrees1/test`.
- 

## RESTORE DESTINATIONS

By default, ONTAP Snapshot for the Quick Recovery Agent restores a volume or file(s) to the client from which it originated; this is referred to as an in-place restore. If desired, you can also restore the data to a different client or file server. Keep in mind the following considerations when performing such restores:

- The destination client must reside in the same CommCell as the client whose data was replicated.
- Note that when you perform restores other than in-place restores, the restored data assumes the rights (i.e., permissions) of the parent directory.

The following section enumerates the types of restore destinations that are supported by ONTAP Snapshot for the Quick Recovery Agent. See [Restore/Recover/Retrieve Destinations - Support](#) for a list of Agents supporting each restore destination type.

---

### IN-PLACE RESTORE

- Same path/destination

---

### OUT-OF-PLACE RESTORE

- Same path/destination
- Different path/destination

---

### **CROSS-PLATFORM RESTORES**

- Same Operating System - Different Version
- Different Operating System

Cross-Platform, File Level restores are supported as follows:

- Windows to Unix
- Unix to Windows
- ONTAP Snapshot to Unix
- ONTAP Snapshot to Windows

Cross-Platform, Volume Level restores are not supported.

---

### **RESTORE TO NETWORK DRIVE/NFS-MOUNTED FILE SYSTEM**

For File Level restores only, you can also Restore to a Network Drive/NFS-Mounted File System. This is similar to an out-of-place restore except that:

- The restored data passes through the client computer to the mapped NFS mount.
- The computer that hosts the share or mounted file system need not be another client within the CommCell.

[Back to Top](#)

---

# Subclients - Quick Recovery Agent

Topics | How To | Related Topics

This feature/product/platform is on Extended Support in this release. See [Deprecated Features, Products, and Platforms](#) for more information.

Overview

Configurable Properties

Things to Consider when Creating and Configuring Quick Recovery Agent Subclients

- All QR Agents
- QR Agent on Windows
- QR Agent on Unix

## OVERVIEW

The following table shows subclient creation and configuration details specific to the Quick Recovery Agent.

| AGENT    | Type of Data      | Default Subclient created during install of the Agent | Supports Default Subclient | Supports User Defined Subclient | Contents of the default subclient when user-defined subclient is present | Other Types of subclients supported by the Agent | Notes |
|----------|-------------------|-------------------------------------------------------|----------------------------|---------------------------------|--------------------------------------------------------------------------|--------------------------------------------------|-------|
| QR Agent | volume level data | No                                                    | No                         | Yes                             | N/A                                                                      | None                                             | None  |

There are several reasons to distribute the storage content across multiple subclients:

- To ensure that all volumes in the subclient are snapped and copied at the same point-in-time. This is important to preserve the consistency of application objects that may span volumes (e.g., SQL Databases or Exchange Storage Groups).
- To schedule subclient operations (e.g. QR Volume Creation/Update) at different times.
- To run jobs on multiple subclients in parallel.
- To associate subclients with different QR Policies.
- For administrative convenience.

Note that you cannot run multiple, concurrent jobs on the same subclient. Distribute the content to multiple subclients to achieve greater parallelism.

## OVERLAPPING CONTENT

QR Agent subclients may have overlapping content. In this scenario, the user creates a QR Volume of the source data at an early point-in-time to one destination volume, while another QR Volume is incrementally updated to a second destination volume. Each of these operations is managed through its own subclient configured with the source volume as content.

When subclients have the same source volume as content, only one subclient is allowed to create incremental updates. Other subclients must select the **Disable incremental update of QR Volumes** option during subclient creation.

When subclients have overlapping content, they must use the same snapshot engine. If the QR policies associated with the subclients are configured to use different snap engines, they must be reconfigured to use the same snap engine.

## CONFIGURABLE PROPERTIES

Once installed, the agent is configured and is therefore able to manage the data or volumes on the client computer. However, you can change certain aspects of the subclient configuration to manage the data in the manner that best suits your needs.

You can view or change the subclient configuration from the Subclient Properties dialog box. The following information can be configured.

## ACTIVITY CONTROL

You can enable or disable all operations for this CommCell object and all objects below it. For more information, see [Activity Control](#).

---

## CONTENT/DATABASES

You can define the content of the subclient. Most agents include a configure button that displays a dialog where you can add or modify the data included as subclient content. For step-by-step instructions, see [Configure Subclient Content](#).

---

## OVERLAPPING SUBCLIENT CONTENT

To configure subclients for overlapping content, select the **Disable Incremental update of QR Volumes** field. For more information, see [Overlapping Content](#).

---

## PRE/POST PROCESSES

You can add, modify or view Pre/Post processes for the subclient. These are batch files or shell scripts that you can run before or after certain job phases. For more information, see [Pre/Post Processes](#).

---

## QR POLICY ASSOCIATION

You can associate a QR Policy to the subclient for QR volume creation operations. For more information, see [QR Policies](#).

---

## SOFTWARE COMPRESSION

You can specify whether software compression is enabled or disabled for the subclient when the selected QR Policy is configured to use a LAN copy manager. For more information, see [Data Compression](#).

---

## NETWORK THROUGHPUT

You can control the network throughput using Network Bandwidth Throttling.

---

## SUBCLIENT NAME

You can rename a subclient. For step-by-step instructions, see [Rename a Subclient](#).

---

## GENERIC SNAP SCRIPTS

This tab is enabled if you had chosen Generic snapshot as your Snapshot Engine Type when you created the QR Policy for the subclient. In this tab, specify the path names for the snap, unsnap, and recovery (or copyback) scripts, and select the host for the scripts.

---

## USER ACCOUNTS

You can define an account with permissions to execute Pre/Post commands for the agent's archive, backup, or volume creation jobs.

See [Quick Recovery Agents: Other User Accounts in User Accounts and Passwords](#) for more information.

---

## THINGS TO CONSIDER WHEN CREATING AND CONFIGURING QUICK RECOVERY AGENT SUBCLIENTS

When creating and configuring subclients for the Quick Recovery Agent, keep in mind the following considerations:

---

### ALL QR AGENTS

- If you add content to a subclient that has already been scheduled and has been performing QR Volume creation and/or updates, the new content will not be included in the current schedule(s). Delete the schedule and create a new schedule to ensure that the new contents are included.
- If a subclient contains multiple volumes, the volumes will be snapped and copied sequentially; however, if you have sufficient hardware resources for simultaneous copies, you can schedule multiple subclients to be copied at the same time. Therefore, in some cases, it may be possible to accelerate your QR Volume creation by splitting a single subclient containing multiple volumes into smaller subclients containing a single volume. Note that application volumes should never be split into multiple subclients; use caution when configuring your subclients for optimal performance.
- The volume must correspond to a physical disk or RAID array.
- Subclients may have overlapping content; however, if two or more subclients overlap, they all must use the same snapshot engine. If the QR policies associated with the subclients are configured to use different snap engines, they must be reconfigured to use the same snap engine in this scenario.
- If you are creating a subclient that will use Generic Enabler snapshot scripts, see [Generic Enablers](#).
- You **cannot** modify subclient content if an incremental backup schedule exists. Modifying subclient content includes adding or deleting a volume from the subclient or changing the subclient's QR policy. These tasks are not allowed if an incremental backup is scheduled or active.

---

### QR AGENT ON WINDOWS

- When configuring Exchange Database volumes as subclient content, see [Creating a QR Volume for use with Exchange](#).
- When configuring Oracle Database volumes as subclient content, see [Creating a QR Volume for use with Oracle Database Instances](#).
- When configuring SQL Database volumes as subclient content, see [Creating a QR Volume for use with SQL Servers](#).
- SQL volumes will not be added to subclient contents if any databases are in the *suspect* or *loading* state, in which case they may not be discovered by the Quick Recovery Agent. Use the SQL Enterprise Manager to change the mode of the database(s).
- QR does not support data protection for applications and application data residing on the C: drive.
- If you are creating a subclient that will use a QR™ Policy with QSnap as the snapshot enabler and Recovery Points enabled, you must specify a cache volume (unless a cache volume had been specified in Client Properties/Advanced). The cache volume specified during Quick Recovery® Agent subclient creation will be used as the cache volume for all snapshots (source and destination) and for all Agents for that particular machine.

---

## QR AGENT ON UNIX

- When configuring Oracle Database volumes as subclient content, see [Creating a QR Volume for use with Oracle Database Instances](#).
- When configuring a subclient to contain a slice 0 volume, the destination volume must also be a slice 0 volume. To specify the slice 0 volume to be used as the destination volume, manually select the destination volume in the QR Volume Creation Advanced Options dialog box.
- If the slice 0 partition is empty, and if the disk is allocated space from slice 1 onwards, then slice 1 will be considered as slice 0 by the agent. Therefore, from whatever slice partition you start to allocate space on the disk, that particular slice will be considered as slice 0.
- A volume created by volume management software other than VxVM is not valid subclient content.

[Back to Top](#)

---

# Subclients - Quick Recovery Agent - How To

[Topics](#) | [How To](#) | [Related Topics](#)

---

[Associate a Subclient to a QR Policy](#)

[Change Account for Executing Pre/Post Commands \(Data Protection\) \(QR Agent on Windows\)](#)

[Configure a Subclient for Pre/Post Processing of Data Protection Operations](#)

[Configure Subclient Content](#)

[Create a New Subclient](#)

[Configure Subclients for Overlapping Content](#)

[Delete a User-Defined Subclient](#)

[Enable or Disable Operations](#)

[Enable Software Compression and Network Bandwidth for a QR Subclient](#)

[Remove a Process from Pre/Post Processing of Data Protection Operations](#)

[Rename a Subclient](#)

[View Subclient Content](#)

---

## ASSOCIATE A SUBCLIENT TO A QR POLICY

### Related Topics

- [QR Policies](#)
- [Subclients - Quick Recovery Agent](#)
- [Generic Enablers](#)

Whenever you create a subclient for the QR Agent, you must associate that subclient to a QR Policy. The QR Policy defines which snap engine and copy manager will be used for operations on the subclient.

### Before You Begin

- Do not change the QR Policy association while a QR Volume Creation operation is running on the subclient.

*Required Capability:* See [Capabilities and Permitted Actions](#)

▶ To associate a subclient to a QR Policy:

1. From the CommCell Browser, right-click the subclient of a QR Agent whose associated QR Policy you want to change, then click **Properties** from the

shortcut menu.

2. Click the General tab of the Subclient Properties dialog box.
3. From the **QR Policy** list, select a QR Policy to associate with this subclient.
4. Click **OK** to save your changes and close the Subclient Properties dialog box.

## CHANGE ACCOUNT FOR EXECUTING PRE/POST COMMANDS (DATA PROTECTION)

*Required Capability:* See Capabilities and Permitted Actions

▶ To change a user account for executing pre/post commands for Data Protection jobs:

1. From the CommCell Browser, expand the tree to view the appropriate level icon for the affected agent.
  - From the agent, instance/partition, or backup set/archive set level, right-click the appropriate icon, click **All Tasks**, and click **New Subclient** from the short-cut menu.
  - From the subclient level, right-click the subclient icon and click **Properties** from the short-cut menu.
2. From the Subclient Properties dialog box, create and/or configure the subclient as appropriate. Then click the **Pre/Post Process** tab.
3. From the **Pre/Post Process** tab, click **Change**.
4. From the User Account dialog box, select one of the account options. If you select **Impersonate User**, type the appropriate user name and password.
5. Click **OK** to save the settings.

## CONFIGURE A SUBCLIENT FOR PRE/POST PROCESSING OF DATA PROTECTION/ARCHIVE OPERATIONS

### Before You Begin

- We recommend not configuring a pre/post process for a subclient that is currently running a data protection or archive operation.
- Verify that there are no pre/post processes already assigned for the subclient.
- Review the Overview and Agent-Specific Guidelines for your agent before configuring pre/post processes for data protection/archive operations.
- Pre-process commands for the *iDataAgents* will be executed only when the necessary resources (e.g., media, library, drive, etc.) are available.

*Required Capability:* Capabilities and Permitted Actions

▶ To configure a subclient for Pre/Post processing of data protection/archive operations:

1. From the CommCell Browser, right-click the subclient for which you want to configure a pre/post process, and then click **Properties** from the shortcut menu.
2. Click the Pre/Post Process tab of the Properties dialog box.
3. For an agent other than the Oracle RAC *iDataAgent*, click inside the space that corresponds to one of the following phases and type the full path of the process that you want executed during that phase. Alternatively, click **Browse** to locate the process (applicable only for paths that do not contain any spaces). For the Oracle RAC *iDataAgent*, click **Browse** for the corresponding process, click the name of the control node client in the Select Client for Browse dialog box, and click **OK**. Then browse for and click the process.
  - PreBackup
  - PreScan
  - PreArchive
  - PreCopy
  - PreSnap
  - PostBackup
  - PostScan
  - PostArchive
  - PostCopy
  - PostSnap
 Click **OK**.
4. If you want to run a Post Process for all attempts to run that job phase, then select the corresponding checkbox.
5. For subclients on Windows platforms, if **Run As** displays **Not Selected**, or if you want to change the account that has permission to run these commands, click **Change**.
  - a. In the User Account dialog box, select **Use Local System Account**, or select **Impersonate User** and enter a user name and password. Click **OK**.

- b. If you selected Local System Account, click **OK** to the message advising you that commands using this account have rights to access all data on the client computer.
6. Click **OK** to save your changes and close the Pre/Post Process tab of the Properties dialog box.

## CONFIGURE SUBCLIENT CONTENT

### Before You Begin

- Review Subclients.
- Do not configure the content of a subclient while the parent node or any sibling subclient has a data protection or archive operation currently running on it.
- Exchange Mailbox *iDataAgents* and Exchange Mailbox/Public Folder Archiver Agents: If you change the contents of the default backup set or archive set then the auto-discover feature will be disabled. If you disable the auto-discovery feature, newly created mailboxes will not be backed up/archived unless they are manually discovered and assigned to a subclient.
- NAS NDMP *iDataAgents*: You must ensure there is no overlap in content between all subclients. Overlap in subclient content will result in loss of data. An existing subclient's contents are not automatically changed when another subclient is added with overlapping contents.
- SharePoint Server *iDataAgent*: The Site Content Database, the Site Collection Database, the Site Database, and the Site Index for the virtual server must all be assigned to the same subclient.
- Lotus Notes Document *iDataAgent*: Review Assigning Restore View Names to Newly-discovered Databases
- QR Agent: Follow these guidelines when adding a volume to a QR Agent subclient:
  - The volume must correspond to a physical disk or RAID array.
  - A volume created by volume management software other than VxVM is not valid subclient content.
  - Subclients may have overlapping content; however, if two or more subclients overlap, they all must use the same snapshot engine. If the QR policies associated with the subclients are configured to use different snap engines, they must be reconfigured to use the same snap engine in this scenario.
- **Caution Against Re-configuring Default Subclient Content**

We recommend that you do not re-configure the content of a default subclient because this would disable its capability to serve as "catch-all" entity for client data. As a result, the likelihood that some data will not get backed up or scanned for archiving would increase.

*Required Capability:* See Capabilities and Permitted Actions

▶ To configure subclient content:

1. From the CommCell Browser, right-click the subclient for which you want to configure content, click **All Tasks** (if applicable) and then click **Properties**.
2. Follow the procedure below that is applicable for your agent:
  - For File System, Active Directory, File Archiver, Exchange Public Folder *iDataAgents*, NDS, and SharePoint Server *iDataAgents* click the Subclient Properties (Content) tab and configure content for the subclient as described below for your agent:
    - For File System, Active Directory, File Archiver, NDS, and SharePoint Server *iDataAgents*: Type the full path of the data that you want to include as subclient content in the **Enter New Content** field, then click **Add**. Optionally, click **Browse** to enter the content. When browsing content while configuring SharePoint subclients, you can add content via multiple selections with the CTRL or SHIFT keys. For Windows, when specifying a UNC Path, click **As User**, and enter the user account information for the domain user with permissions for that path. For NetWare/DNS, see the Notes section below for content path examples. For Unix File Systems, you can enter the mount point of an NFS-mounted file system, see the Notes section below for examples.
    - For Exchange Public Folder *iDataAgents*: Click **Browse**, select folders to include as content, then click **Add**.
    - For the Unix File System *iDataAgents*, to facilitate the management of resource fork data in Apple double-encoded Macintosh files, click **Enable Apple Double Support**.
    - For the Unix File System *iDataAgents*, to view the actual data path for any symbolic link in the subclient content, click **Expand symbolic links of subclient content** and then click **Discover**.
  - For NAS NDMP *iDataAgents*, configure the **Backup Content Path** field(s) as described below, then click **Add**:
    - Click the drop-down list arrow to display the root volumes on the file server. To change the root volume, click one in the list. If you want to refine the content path further, use the space to the right of (or below) the root volume list to enter additional path information. Note the following:
      - For NetApp, the root volume is the mount path of each volume.  
Example: for volume FS1 the root volume will be /vol/FS1.
      - For EMC Celerra, the root volume is the mount point created for a volume.  
Example: for volume FS1 with mount point /FS1 the root volume will be /FS1.
      - For Hitachi, no root volumes are shown in the drop down list. Type the full path of the root volume.  
Example: for volume FS1 with mount point /mnt/FS1 the root volume will be /mnt/FS1.
      - For BlueArc, the root volume is a combination of a descriptor of the path and the volume name.  
Example: for volume FS1 with a mount point of / the root volume will be /\_\_VOLUME\_\_/FS1.



- Optionally, for NetApp NAS NDMP, click **Browse** to enter the content.
  - For Exchange Mailbox and Exchange Mailbox/Public Folder Archiver Agents follow the procedure to Discover and Assign New Mailboxes or Assign Mailboxes to Another Subclient.
  - For Lotus Notes Database and Document iDataAgents follow the procedure to Discover and Assign New Databases or Assign Databases to a Subclient.
  - For DB2, DB2 DPF, Exchange Database, Novell GroupWise, SharePoint Server, SQL Server Database, Sybase, and MySQL iDataAgents, click the Subclient Properties (Content) tab and configure content for the subclient as described below for your agent:
    - For the DB2 iDataAgent, specify whether you want to include the entire database or a subset of this data as content for the subclient. For the DB2 DPF iDataAgent, specify whether you want to include all the affected database partitions or a subset of this data as content for the subclient.
    - For Exchange and GroupWise iDataAgents: Click **Configure**. From the Add/Modify Subclients dialog box click the subclient entry for the database element/Storage Group that you want to add to the new subclient and select the name of the destination subclient from the list that appears. Alternatively, you can select and assign a range of databases/storage groups using the **Change all selected databases/storage groups to list**. Note that you must have at least one database element/Storage Group assigned to this subclient in order to save the configuration.
 

A database/Storage Group that is not configured for a subclient does not appear in the list. This can be the case if the subclient containing the database/Storage Group was deleted. If this happens, click **Discover** to display all databases/Storage Groups.
    - For the SharePoint Server iDataAgent, follow the procedure to Discover and Assign New Data Types.
    - For the Sybase iDataAgent, follow the procedure to Manually Discover Databases.
    - For the MySQL iDataAgent, follow the procedure to Configure MySQL Databases.
  - For the Informix iDataAgent, click the Subclient Properties (Content) tab and define the contents of the subclient. Specifically, establish the backup mode for the data to be backed up, set the backup level, and decide whether to back up the emergency boot file and/or the ONCONFIG file.
  - For the Oracle, SAP for Oracle, or Oracle RAC iDataAgent, click the Subclient Properties (Content) tab and define the contents of the subclient. To configure this subclient for specialized types of backups, follow the appropriate procedure below:
    - Create Subclient for Backing Up Archived Redo Log Files
    - Create Subclient for Backing Up Offline Databases
    - Create Subclient for Backing Up Online Databases
    - Create Subclient for Performing Selective Online Full Backups
  - For SAN iDataAgents, click the Subclient Properties (Content) tab and configure content for the subclient as described below for your agent:
    - Image Level on Unix iDataAgent: Click **Add**. From the Add Content Path dialog box, select the volume(s) that you want to back up (use CTRL + click to select multiple volumes). Click **OK**. The selected volumes are added to the **Contents of subclient** list. These volumes are automatically configured to be CXBF devices. Alternatively, use Volume Explorer per specific scenarios to configure CXBF devices.
 

To configure an unmounted block device or raw device as content, first use Volume Explorer to configure the device as a CXBF device. Then select the configured CXBF device as subclient content. You can ignore the warning that is displayed.

For more information, see When to Use Volume Explorer. For a step-by-step procedure, see Configure a CXBF Device in Volume Explorer.
    - Image Level and Image Level ProxyHost on Windows iDataAgents: Click **Add**. Then in the **Add Content** dialog box, type the full path of the volume or mount point that you want to include as subclient content, then click **Add**. Optionally, click **Browse** to select the content. Click **OK**. The volume or mount point is added to the **Contents of subclient** list. Add additional content by repeating this step.
    - ProxyHost iDataAgent: Select a backup host from the **Backup Host** list. This is the computer to which the BCV is connected. Click **Add**. In the **Content** field of the Add/Edit Content for Subclient dialog box, type the primary host path of the content that you want to back up, or click **Browse** to find and select this data. In the **Backup Host BCV Path** field of the Add/Edit Content for Subclient dialog box, type the path through which the backup host accesses this data on the BCV, or click **Browse** to find and select this path. Click **OK**. The primary host data path and corresponding backup host BCV path are added as a single entry in the **Contents of subclient** list. To add additional entries, repeat these steps. Refer to Notes below for more information.
  - For Quick Recovery Agents, click the Subclient Properties (Content) tab and configure the following options:
    - Click **Add Volume**. From the Adding Volume dialog box, select volume(s) that you want to add to the subclient content (use CTRL + click to select multiple volumes). You can add/edit additional advanced options for the selected volume by select **Advanced** on the Adding Volume dialog box. Click **OK**.
    - Click Add App to select an application and associated volumes. Click **OK**.
 

Any instances you intend to protect and recover with the QR Agent must be configured in the QR Agent properties Authentication tab. They will not appear in the Add App dialog box if they are not configured. Only volumes containing datafiles and archive log files will be detected by Add App. Volumes containing control files and redo log files will not be detected.

For a clustered Exchange Server, if you are *not* using VSS to perform an online quiesce, sufficient permissions are required in order to be able to perform an offline quiesce; in such cases, ensure that the **User Name** specified has Exchange Administrator rights.

See also Configure Subclients for Overlapping Content.
3. Click **OK** to save your content configuration.

## NOTES

- Content examples for NetWare are **OU=prospects.O=engineering.[Root]**, (for NDS content), and **SYS:\public** (for File System content).
- Content examples for adding an NFS-mounted file system to subclient content of a Unix File System *iDataAgent*:
  - `/mountpointA` to include the entire file system at mountpointA
  - `/mountpointA/projects` for only the *projects* directory within the file system at mountpointA.
- Informix subclients include one or more dbspaces. As databases are added to the dbspaces, the subclients are updated automatically.
- Exchange Mailbox *iDataAgents* and Exchange Mailbox/Public Folder Archiver Agents: Initially, all unconfigured mailboxes are assigned to the default subclient. You can create a new subclient and reassign mailboxes to this new subclient (within the same backup set/archive set). Once assigned, the mailboxes become part of the content of the new subclient.
- SharePoint Server *iDataAgent*: Initially, all unconfigured data types are assigned to the default subclient. You can create a new subclient and reassign data types to this new subclient. Once assigned, they become part of the content of the new subclient.
- ProxyHost *iDataAgent*: The primary host data path is backed up by the subclient and is the path through which the backup host accesses this data on the BCV. A primary host path and its corresponding backup host path are listed in the following format:

**<primary\_host\_path> --> <backup\_host\_path>**

For example, assume that you want to back up the **D:\data** directory from your primary host and **D:\** is mirrored by a BCV, which is mapped to the backup host as **F:\**. Consequently, the path to this data on the backup host is **F:\data**. When you add this directory to a subclient, it is listed in the **Contents of subclient** pane as **D:\data --> F:\data**.

The primary host path in the **Content** field is used for browse and restore purposes. However, it is the data in the **Backup Host BCV Path** which is actually backed up. If these two paths do not accurately correspond, the path that appears when data is browsed for restore does not accurately reflect the data that will be restored. In the example given above, assume that **D:\data** is entered in the **Content** field, while **F:\data1** is accidentally entered in the **Backup Host BCV Path**. If you browse and select **D:\data** to be restored, it is actually **D:\data1** that is restored. (Remember, **F:\Data1** is the path on the backup host that corresponds to **D:\data1** on the primary host.)

## CREATE A NEW SUBCLIENT

### Before You Begin

- Review Subclients.
- Do not create a subclient while the parent node or any sibling subclient has a data protection or archive operation currently running on it.
- In cases where a new subclient is created with the same name as a deleted subclient, the system will append a Unix time stamp to the deleted subclient's name in data protection job history reports and views to distinguish the two subclients. For example, `subclientname_1104257351`.
- Informix *iDataAgents*: If you will be using the Informix ONBAR utility to create backup and restore scripts, you need not create subclients. Otherwise, if you will be using the CommCell Console to back up and restore Informix database objects (subsets/dbspaces), then you will need to create a subclient.
- ProxyHost *iDataAgents*: If you are using a BCV, you must prepare a batch file or a shell script file on the backup host containing commands to synchronize and split the BCV. The Resource Pack includes information on configurations for these batch files or shell scripts, as well as examples that apply to specific applications and hardware (e.g., Exchange databases in an EMC Symmetrix environment). See Resource Pack for more information on the Resource Pack. The ProxyHost *iDataAgent* also requires that you set permissions for the batch/shell script file on the backup host.
- SQL Server Database *iDataAgents*: When running on Windows Server 2003 and VSS is enabled, the **New Subclient** command is not available.
- PostgreSQL *iDataAgents*: Once you configure the PostgreSQL instance, the system automatically generates the default backup sets and default subclients. However, you can use the CommCell Console to create user-defined subclients for dump backup sets to distribute some of the database content. You cannot create user-defined subclients for FS backup sets.

*Required Capability:* See Capabilities and Permitted Actions

▶ To create a new subclient:

1. From the CommCell Browser, right-click the node (agent/backup set/archive set/instance) for which you want to create a new subclient, click **All Tasks** (if applicable), and then simply click **New Subclient** for most agents.
  - For the SQL Server *iDataAgent*, expand **New Subclient** and click either **Database** to include individual databases or **File/File Group** to include database elements.
2. Click the General tab or General (Quick Recovery Agent) tab of the Subclient Properties dialog box and type the name (up to 32 characters) of the subclient that you want to create.
  - For supported agents identified in Support Information - Snapshot Engines, you can select a QSnap option to snap data and then perform a data protection operation on the data.
  - For Image Level on Unix and Image Level ProxyHost on Unix, use the **Incremental Support Using** field to configure either a CXBF subclient or a checksum subclient and to enable incremental support for either subclient type.
  - For QR Agents, you must also select a QR Policy from the **QR Policy** list.
  - For the Windows *iDataAgents* that support VSS, you can optionally Enable VSS on a Subclient.

3. Select other options from the General tab as appropriate for the agent.
4. Click the **Content** or **Databases** tab of the Subclient Properties dialog box and Configure Subclient Content as appropriate for your agent.
5. For all agents (except QR), click the Storage Device (Data Storage Policy) tab of the Subclient Properties dialog box, then select a data storage policy to associate with this subclient from the storage policy list.
  - o For the DB2 and DB2 DPF *iDataAgents*, you can also change the number of data backup streams. For the DB2 DPF *iDataAgent*, the default stream threshold should be equal to the total number of database partitions for the subclient.
  - o For SQL Server *iDataAgents*, you can also click the Storage Device (Log Storage Policy) tab of the Subclient Properties dialog box, then select a log storage policy to associate with this subclient from the storage policy list and select the number of backup streams for transaction log backup jobs.
  - o For 1-Touch for Unix, it is strongly recommended that the storage policy that you select for the subclient configured for 1-Touch use a MediaAgent on a different computer. If you do this, and if the system crashes, the media will not have to be exported to another MediaAgent in order to recover the system.
6. For Oracle and DB2 *iDataAgents*, click the Backup Arguments (Oracle) or Backup Arguments (DB2, DB2 DPF) tab of the Subclient Properties dialog box and Configure Backup Arguments as appropriate for your agent. Note that the backup arguments for Informix are located on the Content tab.
7. For Migration Archiver Agents, click the **Archiving Rules** or **Rules** tab of the Subclient Properties dialog box and configure archiving rules as appropriate for your agent. In order to perform rules-based migration archiving operations, the **Disable All Rules** checkbox must be cleared.
 

If the File Archiver for Windows supports Data Classification, several filter-like configuration fields are defined as archiving rules and are available from the Subclient Properties (Rules) tab. If you want to define content and archiving rules based on file attributes other than volumes, size, and modified time (i.e., if you want to customize your rules), click the Advanced tab and configure as appropriate. Also, stub management options can be configured from the Stub Rule tab. See Configure Archiving Rules - File Archiver Agents for step-by-step instructions.
8. For ProxyHost and Image Level ProxyHost *iDataAgents*, click the Pre/Post Process tab of the Subclient Properties dialog box. In the **PreScan** field, type the path to the batch file/shell script file that contains those commands that are to run before each backup of the subclient, or click **Browse** to locate and select this file. For ProxyHost and Image Level ProxyHost, the file must reside on the backup host or primary host.
9. Optionally (if supported for your agent) you can:
  - Add a Data Protection or Discovery Filter for a Subclient on the Filters tab.
  - Configure a Subclient for Pre/Post Processing of Data Protection/Archive Operations on the Pre/Post Process tab.
  - Enable Software Compression for a Subclient on the Software Compression tab of the **Storage Device** tab.
  - Configure the Subclient for Data Encryption on the Encryption tab.
  - Enable or Disable Operations for this subclient on the Activity Control tab.
  - Configure Mailbox Stores for Auto-Discovery on the Auto-discovery tab.
  - Configure the Subclient for 1-Touch on the 1-Touch Recovery tab.
  - View or change the user group security associations for this subclient from the Security tab.
  - Determine location from where archive logs will be backed up or deleted from the Log Destinations tab.
10. Click **OK** to save the subclient configuration. For QR Agents, this procedure is now complete. For all other agents, continue on to the next step.
11. The Backup Schedule dialog box advises you to schedule data protection operations for your new subclient. It is recommended you elect to set a schedule now. You can also associate this subclient with an All Agent Types schedule policy (which is automatically created by the system, or can be a user defined Data Protection schedule policy). If you have already associated a schedule policy at a previous level (Backup Set/Instance, Agent, Client, or Client Computer Group) the schedules defined in the Schedule Policy will be automatically applied to the new subclient. See Schedule Policy for more information.
  - o If you want to associate this subclient with an All Agent Types schedule policy, click **Associate with Generic Schedule Policy**, and then select that schedule policy from the drop-down list box. Click **OK**.
  - o If you want to associate this subclient with a specific schedule policy, click **Associate to schedule policy**, and then select the schedule policy from the drop-down list box. Click **OK**.
  - o If you have selected to define a schedule for this subclient:
    - Click **Schedule**.
    - From the Backup/Archive Options dialog box, select the type of data protection operation that you want to schedule.
    - If you want to set Advanced Backup/Archive Options, click **Advanced**.
    - After selecting the data protection type and any advanced options, click **OK**. The **Schedule Details** dialog box appears.
    - From the Schedule Details tab, select the scheduling options that you want to apply, then click **OK**.
  - o If you don't want to create a data protection schedule at this time, click **Do Not Schedule**, and then click **OK**.

This task is now complete.

---

## CONFIGURE SUBCLIENTS FOR OVERLAPPING CONTENT

**Related Topics:**

- Subclients - Quick Recovery Agent

**Before You Begin**

- Review Overlapping Content.
- Subclient overlap is intended for use when you would like to have the same source volume have both a full QR volume to hold a copy of your data from a healthy point-in-time, and an incrementally updated QR volume to hold a more recent copy of your data. Below is a general outline of the steps necessary to add the same volume as content to two different subclients.

*Required Capability:* See Capabilities and Permitted Actions

▶ To configure subclients for overlapping content:

1. Create a new subclient as outlined in Create a New Subclient and add the volume as contents.
2. Create a new subclient as outlined in Create a New Subclient, add the volume as contents and assign the subclient to a different QR Policy than the one assigned to the first subclient you created with this volume as content. Note that in order for subclient overlap to work, both QR policies must use the same snap engine.
3. When a volume has been added as content to more than one subclient, the system will prompt you to disable incremental updates on one of the subclients. Click **OK**. From the Subclient Properties (General) tab, select the **Disable Incremental update of QR Volumes** option. The subclient for which you select this option will only be able to perform full QR volume creations. The other subclient that does not have the **Disable Incremental update of QR Volumes** option selected is able to perform both full QR volume creations and incremental updates.
4. The subclients are now setup for overlapping content. When scheduling the subclients, note that QR operations on the two subclients will not execute in parallel (i.e., the first subclient cannot perform an incremental update of the volume while the other subclient is performing a full QR Volume creation).

## DELETE A USER-DEFINED SUBCLIENT

**Related Topics:**

- Command Line Interface - qdelete subclient
- Subclients

*Required Capability:* See Capabilities and Permitted Actions

▶ To delete a user-defined subclient:

1. From the CommCell Browser, right-click the user-defined subclient that you want to delete, and then click **Delete** from the shortcut menu.
2. A confirmation message is displayed, asking if you want to delete the subclient. Click **No** to cancel the deletion and retain the subclient, or click **Yes** to continue the deletion. If you click **Yes**:
  - the subclient, and any data that may have been protected/archived by the subclient are logically deleted, and you can no longer access the corresponding data for recovery/retrieve purposes. However, the data remains valid for the length of time specified by the associated retention period. Some agents allow you to browse data from a deleted subclient provided that the Browse Data Before date and time precedes the time that the user-defined subclient was deleted.
  - for agents that support a default subclient, once the user-defined subclient is deleted its contents are automatically reallocated to the default subclient the next time a data protection/archive/discovery operation is run on the default subclient to ensure data protection coverage.
  - the system deletes the selected subclient node and removes it from the CommCell Browser.
  - the system deletes any data protection/archive and recovery/retrieve job schedules that are associated with the subclient.

## ENABLE OR DISABLE OPERATIONS

*Required Capability:* See Capabilities and Permitted Actions

| Level                 | Capability                                                             |
|-----------------------|------------------------------------------------------------------------|
| CommCell              | Administrative Management with CommCell level association              |
| Client Computer Group | Administrative Management with Client Computer Group level association |
| Client                | Agent Management with Client level association                         |
| Agent                 | Agent Management with Agent level association                          |
| Subclient             | Agent Management with Subclient level association                      |

▶ To enable or disable activity control at the CommCell, client computer group, client, agent, or subclient levels:

1. From the CommCell Browser, right-click the CommServe, client computer group, client computer, agent, or subclient, and then click **Properties** from the short-cut menu.

2. From the Activity Control tab of the associated Properties dialog box, select or clear option(s), as desired.
3. Click **OK**.



Disabled data management and/or data recovery operations are displayed with client and/or agent icon changes in the CommCell Browser. For a comprehensive list of all icons in the CommCell Console, see CommCell Console Icons.

## ENABLE OR DISABLE SOFTWARE COMPRESSION AND NETWORK BANDWIDTH FOR A QR SUBCLIENT

### Before you Begin

- Do not change the software compression while a QR Volume Creation operation is running on the subclient.

*Required Capability:* Capabilities and Permitted Actions

▶ To enable software compression and Network Bandwidth for a QR subclient:

1. From the CommCell Browser, right-click the subclient of a QR Agent for which you wish to enable software compression, then click **Properties** from the shortcut menu.
2. Click the General tab of the Subclient Properties dialog box.
3. Select the **Compression ON** option to enable software compression.
4. Click the **Throttle Network Bandwidth (MB/HR)** option and then enter the throughput as needed.
5. Click **OK** to save your changes.

This task is now complete.

## REMOVE A PROCESS FROM PRE/POST PROCESSING OF DATA PROTECTION/ARCHIVE OPERATIONS

### Before You Begin

- We recommend not removing a pre/post process for a subclient that is currently running a data protection or archive operation.
- Review the Overview and Agent-Specific Guidelines for your agent before removing pre/post processes for data protection/archive operations.

*Required Capability:* Capabilities and Permitted Actions

▶ To remove a process from Pre/Post processing of data protection/archive operations:

1. From the CommCell Browser, right-click the subclient for which you want to remove a pre/post process, and then click **Properties** from the shortcut menu.
2. Click the Pre/Post Process tab of the Subclient Properties dialog box.
3. Click the text inside the space that corresponds to one of the following phases for which you want a pre/post process removed, then press the **Delete** key:
  - PreScan
  - PreArchive
  - PreCopy
  - PreSnap
  - PostBackup
  - PostScan
  - PostArchive
  - PostCopy
  - PostSnap
4. Repeat Step 3 for any additional processes that you want to remove.
5. Click **OK**.

## RENAME A SUBCLIENT

### Before You Begin

- You can rename a subclient at any time. However, we recommend that you don't rename a subclient while a data protection or archive operation is running on that subclient.

- In cases where a subclient is renamed using the same name as a deleted subclient, the system will append a Unix time stamp to the deleted subclient's name in data protection job history reports and views to distinguish the two subclients. For example, `subclientname_1104257351`.

*Required Capability:* See Capabilities and Permitted Actions

▶ To rename a subclient:

1. From the CommCell Browser, right-click the subclient that you want to rename, and then click **Properties** from the shortcut menu.
2. From the Subclient Properties (General) tab, or the QR Agent Subclient Properties (General) tab, type the new name in the **Subclient Name** field, and then click **OK**.

The CommCell Browser updates the subclient with its new name. The new name will also be reflected in any associated schedules and reports.

---

## VIEW SUBCLIENT CONTENT

*Required Capability:* See Capabilities and Permitted Actions

▶ To view content of a subclient:

1. From the CommCell Browser, right-click the subclient whose content you want to view, then click **Properties**.
  2. From the Subclient Properties dialog box, click the **Content** tab (or **Databases** tab for Lotus Notes) to view the contents of the subclient.
  3. Click **OK** to close the dialog box.
- 

[Back To Top](#)

# QR Policies

Topics | How To | Related Topics

This feature/product/platform is on Extended Support in this release. See [Deprecated Features, Products, and Platforms](#) for more information.

Overview

Deleting QR Policies

## OVERVIEW

The QR Policy associated with a subclient determines how the volumes in that subclient will be snapped and copied to QR Volumes. When creating a QR Policy, a snapshot engine and copy manager are selected.

The QR Policy also includes a built-in retention policy that defines how long QR Volumes will be retained after creation. Stale QR Volumes are automatically deleted (and their resources returned) each time the Data Aging administration job runs on the CommServe. To prevent QR Volumes from being pruned, you can specify an infinite retention policy. To prevent an individual QR Volume from being pruned, you can set the In Use flag on the QR Volume Properties screen.



- QR Volumes that are scheduled for incremental updates are never pruned.
- If you are creating a QR Policy for use with ONTAP SnapVault or OSSV SnapVault, review the respective overview first.
- Review Recovery Points before creating a QR Policy for use with Recovery Points.

## DELETING QR POLICIES

When a QR policy is deleted, all the QR Volume information associated with the QR policy will be removed from the CommServe database. Therefore, once the QR policy is deleted, the QR Volumes associated with the QR policy cannot be recovered. Note that a QR Policy cannot be deleted if there are any Data Aging or recovery jobs running on this QR Policy.

A QR Policy cannot be deleted if it is associated with a subclient. If you want to delete a QR Policy that is associated with a subclient, then you must do the following:

- Delete any currently associated subclients or
- Reassign the existing subclients to another QR Policy.

You may decide to delete a QR policy if:

- You determine that you do not need the QR Volumes that were created through that QR policy.
- QR Volumes no longer exist on the QR policy and you have no plans to use the policy for future QR Volume creations.

[Back to Top](#)

# QR Policies - How To

Topics | How To | Related Topics

[Create a QR Policy](#)

[Delete a QR Policy](#)

[Modify QR Policy Properties](#)

## CREATE A QR POLICY

*Required Capability:* Capabilities and Permitted Actions

The system allows you to create a QR policy in the CommCell Console from the QR Policies level.

▶ To create a QR policy:

1. From the CommCell Browser, right-click the **QR Policies** icon, and then click **Create New QR Policy** from the short-cut menu.

2. Type the policy name (Up to 32 characters).
3. Select the **Snapshot Engine Type** from the dropdown menu.



- If you are creating snapshot scripts using the Generic Enabler feature, select **Generic snapshot** as your Snapshot Engine Type.
- Depending on the Snapshot Engine you select, some options may not be available for configuration.

4. Select the **Retention Policy** by checking Infinite (default), or by entering the number of hours/days/weeks.
5. Select the **Enable QR Volume Creation** check box to activate the **Copy Manager** and **Scratch Pool** menus.



This option is automatically disabled if you selected **Generic snapshot** as your Snapshot Engine Type in Step 3.

6. Select the **Copy Manager** from the list of available Copy Managers that currently exist for this QR Volume. (Refer to LAN Copy Manager.)
7. Select the **Scratch Pool** from the list of available scratch pools that currently exist for this QR Volume.
8. If this QR Policy will be used to create Recovery Points, select **Enable Recovery Points**, and specify the maximum number of Recovery Points that will be retained, up to a maximum of 32.
  - When you select **Enable Recovery Points**, you are prompted to enter a location for the cache partition for the snapshots of the QR Volume. This location cannot be on the QR Volume.
  - When you are satisfied with the QR Policy's configuration, click **OK**. The new QR Policy appears in the CommCell Browser.

## DELETE A QR POLICY

**WARNING!** When a QR policy is deleted, all the QR Volume information associated with the QR policy will be removed from the CommServe database. Therefore, once the QR policy is deleted, the QR Volumes associated with the QR policy cannot be recovered. Note that a QR Policy cannot be deleted if there are any Data Aging or recover jobs running on this QR Policy.

*Required Capability:* Capabilities and Permitted Actions

▶ To delete a QR policy:

1. Select **QR Policies** from the CommCell Browser, and right-click the QR policy that you wish to delete.
2. Click **Delete** from the short-cut menu.

The system prompts you to confirm the deletion of the QR policy.

3. To confirm the deletion, click **Yes**. The deleted QR Policy is removed from the CommCell Browser.

## MODIFY QR POLICY PROPERTIES

It is recommended that you do not change any of the QR policy properties while the QR policy being used by an operation (e.g., QR Volume Creation, QR Volume Recovery, etc.)

*Required Capability:* Capabilities and Permitted Actions

▶ To change the properties of a QR policy:

1. Select **QR Policies** from the CommCell Browser, click **Properties** and right-click the QR policy that you wish to modify.
2. Change any of the following properties:
  - QR policy name (Up to 32 characters).
  - Snapshot Engine Type from the dropdown menu.
  - Retention Policy by checking Infinite (default) or enter the number of hours/days/weeks.
  - **Enable QR Volume Creation** check box to activate the **Copy Manager** and **Scratch Pool** menus.
  - Copy Manager from the list of available Copy Managers that currently exist for this QR Volume.
  - Scratch Pool from the list of available scratch pools that currently exist for this QR Volume.
  - If this QR Policy will be used to create Recovery Points, select **Enable Recovery Points**, and specify the maximum number of Recovery Points that will be retained, up to a maximum of 32.
    - When you select **Enable Recovery Points**, you are prompted to enter a location for the cache partition for the snapshots of the QR Volume. This location cannot be on the QR Volume.
3. Once you have made your changes, click **OK**.



[Back To Top](#)

# Scratch Volume Pools

Topics | How To | Related Topics |

---

This feature/product/platform is on Extended Support in this release. See [Deprecated Features, Products, and Platforms](#) for more information.

Overview

QR Volumes and Scratch Volume Pools

Scratch Volume Pool Considerations

Recovery and Scratch Volume Pools

---

## OVERVIEW

A Scratch Volume Pool is a repository of destination volumes that are available for use by a QR Volume Creation operation. The system does not create Scratch Volume Pools by default. The Pools must be manually created and assigned to a QR Policy. Scratch Volume Pools are created in the Storage Resources section of the CommCell Browser.

---

## QR VOLUMES AND SCRATCH VOLUME POOLS

When creating a QR Volume, you may select a specific destination volume or allow the Quick Recovery Agent to select one automatically. Destination volumes must be selected from the Scratch Volume Pool associated with the QR Policy. The Quick Recovery Agent will select an un-allocated volume from the pool unless a specific volume is chosen by the user. The agent will attempt to allocate the smallest available volume that is of a sufficient size to hold the primary volume's contents. However, it is recommended that you create different Scratch Volume Pools containing volumes of consistent size, and assign the pool of appropriate size to each QR Policy and/or subclient as needed.

A disk volume must be detected and configured by Volume Explorer before it can be assigned to a Scratch Volume Pool. A single volume may only belong to one Scratch Volume Pool at a time. The Quick Recovery Agent keeps track of scratch volume allocations and will prevent multiple QR Volume operations from attempting to use the same scratch volume. If a QR Volume Creation/Update operation fails unexpectedly, it may leave the destination volume in a reserved state. The CommCell Archive Pruning operation will detect and clean up any such volumes the next time it is run.

Only QSnap (CXBF) devices can be used as destination volumes when using QSnap and the Quick Recovery Agent. Otherwise, the destination volume could become inconsistent with the source. If a QSnap (CXBF) device becomes de-configured, it may be chosen from the Scratch Volume Pool as a destination volume. To avoid this, manually choose volumes when creating QR destination Volumes; look for "cxbf" in the Device Name.

---

## SCRATCH VOLUME POOL CONSIDERATIONS

When creating a Scratch Volume Pool, you must ensure that:

- The destination volumes selected for the scratch volume pool have a capacity greater than or equal to the size of the source volume(s) being copied.
- The destination volumes have been configured in Volume Explorer with the correct WWN and LUN for proper detection in the SAN network.
- When using a LAN copy manager configuration, a Scratch Volume Pool should only contain volumes from one client. In a SAN environment, volumes from multiple clients can be added.

The user can manually add and delete volumes to and from an existing pool. Once a copy has been created and data has been copied to a destination volume in your associated scratch volume pool, the volume will be marked **locked**, which can be seen from the Volume Locked check box in the right panel of the scratch volume pool. Once this volume is locked, any subsequent QR Volume creation jobs will not be allowed to use it as a destination volume. This is to protect the volume from being overwritten. Any destination volumes which are locked but no longer in use by any QR Volumes will be released during the next data aging job on the CommServe.

You can create any number of additional scratch volume pools and assign them to any QR policies that you wish. The user can also add volumes from different clients into the same pool, which allows for more flexibility in QR volume creation. The user cannot, however, overlap volumes. Scratch volume pools are mutually exclusive. For instance, the F: drive on client1 cannot be used in more than one scratch volume pool. To prevent an overlap from occurring, available volumes to be added to the scratch volume pool are filtered so that assigned volumes will not be viewable from the list.

---

## RECOVERY AND SCRATCH VOLUME POOLS

When a recover is performed, the mount point or drive letter of the destination volume is switched with that of the source volume. In other words, the scratch volume pool will show the source volume's drive letter in place of the destination volume. This concept can be confusing, so an example is provided below:

If the user performs a QR Volume creation with F: drive (source) to the Z: drive (destination), then a QR recover will unmount the source (F:) and destination volume (Z:) and mount the QR volume as F:.

This is an automatic feature of Recovery and requires no user intervention. From this point the user has some options. The most common choice would be to delete the QR Volume and remove it from the scratch volume pool. The other option is that the user could choose to leave the recovered volume in a locked status.

[Back to Top](#)

---

## Scratch Volume Pools - How To

[Topics](#) | [How To](#) | [Related Topics](#) |

---

[Create a Scratch Volume Pool](#)

[Delete a Scratch Volume Pool](#)

[Modify Scratch Volume Pool Properties](#)

[Remove a Volume from a Scratch Volume Pool](#)

---

### CREATE A SCRATCH VOLUME POOL

*Required Capability:* See Capabilities and Permitted Actions

▶ To create a Scratch Volume Pool:

1. From the CommCell Browser, right click the Scratch Volume Pool, and then click **Create Pool** from the short-cut menu.
2. In the Create Scratch Volume Pool dialog box, enter the name of the new Scratch Volume Pool. If this scratch pool consists of ONTAP volumes for use with SnapVault, select the **Use for SnapVault Only** checkbox. To use the scratch pool for SnapMirror, select the **Use for SnapMirror Only** checkbox.
3. When you are satisfied, click **OK**.

The new Scratch Volume Pool appears in the CommCell Browser.

---

### DELETE A SCRATCH VOLUME POOL

A scratch volume must be empty before it can be deleted. Remove all volumes from the scratch volume pool before deletion.

*Required Capability:* See Capabilities and Permitted Actions

▶ To delete a Scratch Volume Pool:

1. From the CommCell Browser, right-click the Scratch Volume Pool that you want to delete, and then click **Delete**.
2. A confirmation prompt appears, asking if you are sure that you want to delete this Scratch Volume Pool. Click **Yes** to delete.

The Scratch Volume Pool is deleted and the CommCell Browser display is updated.

---

### MODIFY SCRATCH VOLUME POOL PROPERTIES

When using a LAN copy manager configuration, a Scratch Volume Pool should only contain volumes from one client. In a SAN environment, volumes from multiple clients can be added.

*Required Capability:* See Capabilities and Permitted Actions

▶ To modify the properties of a Scratch Volume Pool:

1. From the CommCell Browser, right-click the Scratch Volume Pool whose property you want to change, and then click **Properties** from the short-cut menu.
2. The Scratch Volume Pool Properties (General) tab displays the number of total scratch volumes and available scratch volumes that the Scratch Volume Pool contains. This tab also allows you to change the Scratch Volume Pool's name. Click in the Name field and enter a value to change the selected property.

3. The Scratch Volume Pool Properties (Volumes) tab displays the Volume Locked status, Client, Mount Paths, and capacity of the scratch volumes. Click **Add** to open the Add Volume to Scratch Volume Pool dialog box.
  4. From the Add Volume to Scratch Pool dialog box, select a scratch volume from a list of available scratch volumes, and click **Add**. When you are satisfied with your additions, click **OK**.
  5. When you are satisfied with your changes, click **OK**.
  6. Click **OK** to save your entries.
- 

## REMOVE A VOLUME FROM A SCRATCH VOLUME POOL

When using a LAN copy manager configuration, a Scratch Volume Pool should only contain volumes from one client. In a SAN environment, volumes from multiple clients can be added.

*Required Capability:* See Capabilities and Permitted Actions

▶ To remove a scratch volume from a Scratch Volume Pool:

1. In the left pane of the CommCell Browser, select the Scratch Volume Pool containing the scratch volume that you want to delete. The contents of the Scratch Volume Pool are displayed in the right pane of the Browser.
2. From the right pane of the CommCell Browser, right-click the scratch volume that you want to delete, and then click **Delete**.
3. If you are sure that you want to delete the scratch volume, click **Yes** in the Confirm Delete prompt that appears.

The scratch volume is deleted from the Scratch Volume Pool and is no longer available to the QR Agent for use.

---

[Back To Top](#)

# QR Volume Creation History

Topics | How To | Related Topics

The **QR Volume Creation Job History Filter** dialog box allows you view detailed, historical information about quick recovery volume creation operations for the Quick Recovery Agent.

For information on Job Details displayed in the Job History, see [Viewing Job Information](#).

Once chosen, your filter options are then displayed in the QR Volume Creation Job History of QR Agent window. From this window you can view more detailed information such as the:

- Details of the QR volume creation job.
- Events of the QR volume creation job.
- Log files of the QR volume creation job.



After releasing a QR Volume, only the failed jobs will be retained in the job history.

## QR Volume Creation History - How To

Topics | How To | Related Topics

[View QR Volume Creation Job History for the Quick Recovery Agent](#)

[View Job History Details](#)

[View the Events of a Job History](#)

[Viewing the Log Files of a Job History](#)

### VIEW QR VOLUME CREATION JOB HISTORY FOR THE QUICK RECOVERY AGENT

▶ To view the job history of QR Volume Creation operations:

1. From the CommCell Browser, right-click an agent whose QR Volume Creation history you want to view, click **View**, then select **View Job History -> QR Volume Creation**.
2. From the QR Volume Creation Job History Filter, select the desired options and then click **OK**.
3. The system displays the options you selected in the QR Volume Creation Job History window.
4. Click **OK**.

### VIEW JOB HISTORY DETAILS

*Required Capability:* See [Capabilities and Permitted Actions](#)

▶ To view the details of a job history:

1. From the CommCell Browser, right-click the entity (e.g., subclient, client computer, etc.) whose job history you want to view, click **View**, and then click job history.
2. From the Job History Filter dialog box, select the filter options that you want to apply and click **OK**.
3. From the Data Management Job History window, right-click the job whose job details you want to view, and then click **View Job Details**.
4. The Job Details dialog box appears, displaying detailed job history in General, Details, Phase Details and Attempts tabs for the selected job.
5. Click **OK**.



If viewing the details of a job with a pending or failed status, the **Reason for Job Delay** field will contain an Error Code, which, if clicked, will launch the customer support website displaying troubleshooting article(s) related to the specific issue.

## VIEW THE EVENTS OF A JOB HISTORY

*Required Capability:* See Capabilities and Permitted Actions

▶ To view the events associated with a job:

1. From the CommCell Browser, right-click the entity (e.g., subclient, client computer, etc.) whose job history you want to view, click **View**, and then click **Job History**.
  2. From the Job History Filter dialog box, select the filter options that you want to apply and click **OK**.
  3. From the Data Management Job History window, right-click the job whose job details you want to view, and then click **View Events**.
  4. The All Found Events window gets displayed. If no events were found for the backup, a message is displayed to that effect.
  5. Click **Close**.
- 

## VIEW THE LOG FILES OF A JOB HISTORY

*Required Capability:* See Capabilities and Permitted Actions

▶ To view the log files of a Job History:

1. From the CommCell Browser, right-click the entity whose job history you want to view, and then click to view a job history.
  2. From the job history filter window select the filter options, if any, that you want to apply, and then click **OK**.
  3. From the job history window, right-click the job whose log files you want to view, and then click **View Logs**.
  4. The contents of the log file related to the selected job history are displayed in the **Log File for Job n** window.
- 

[Back to Top](#)

# QR Volume Recovery History

Topics | How To | Related Topics

The *Recovery History Filter for QR Agent* dialog box allows you to view detailed historical information about quick recovery jobs.

For information on Job Details displayed in the Job History, see *Viewing Job Information*.

Once your filter options are chosen, the Recovery Job History window displays the quick recovery jobs that meet the criteria you selected in the Recovery History Filter for QR Agent. From this window you can view more detailed information such as the:

- Details of the QR volume recovery or Copyback job
- Events of the QR volume recovery or Copyback job
- Log files of the QR volume recovery or Copyback job
- Source Volume
- Destination Volume



- The source and destination fields are not supported for Copyback history.
- After releasing a QR Volume, only the failed jobs will be retained in the job history.

## QR Volume Recovery History - How To

Topics | How To | Related Topics

View QR Volume Creation Job History for the Quick Recovery Agent

View Job History Details

View the Events of a Job History

Viewing the Log Files of a Job History

### VIEW QR VOLUME RECOVERY JOB HISTORY FOR THE QUICK RECOVERY AGENT

▶ To view the job history of QR Volume recovery operations:

1. From the CommCell Browser, right-click an agent whose QR Volume recovery history you want to view, click **View**, then click **Job History > QR Volume Recovery**.
2. From the QR Volume Recovery History Filter, select the desired options and then click **OK**.
3. The system displays the options you selected in the QR Volume Recovery Job History window.
4. Click **OK**.

### VIEW JOB HISTORY DETAILS

*Required Capability:* See Capabilities and Permitted Actions

▶ To view the details of a job history:

1. From the CommCell Browser, right-click the entity (e.g., subclient, client computer, etc.) whose job history you want to view, click **View**, and then click job history.
2. From the Job History Filter dialog box, select the filter options that you want to apply and click **OK**.
3. From the Data Management Job History window, right-click the job whose job details you want to view, and then click **View Job Details**.
4. The Job Details dialog box appears, displaying detailed job history in General, Details, Phase Details and Attempts tabs for the selected job.
5. Click **OK**.



If viewing the details of a job with a pending or failed status, the **Reason for Job Delay** field will contain an Error Code, which, if clicked, will launch the customer support website displaying troubleshooting article(s) related to the specific issue.

## VIEW THE EVENTS OF A JOB HISTORY

*Required Capability:* See Capabilities and Permitted Actions

▶ To view the events associated with a job:

1. From the CommCell Browser, right-click the entity (e.g., subclient, client computer, etc.) whose job history you want to view, click **View**, and then click **Job History**.
  2. From the Job History Filter dialog box, select the filter options that you want to apply and click **OK**.
  3. From the Data Management Job History window, right-click the job whose job details you want to view, and then click **View Events**.
  4. The All Found Events window gets displayed. If no events were found for the back up, a message is displayed to that effect.
  5. Click **Close**.
- 

## VIEW THE LOG FILES OF A JOB HISTORY

*Required Capability:* See Capabilities and Permitted Actions

▶ To view the log files of a Job History:

1. From the CommCell Browser, right-click the entity whose job history you want to view, and then click to view a job history.
  2. From the job history filter window select the filter options, if any, that you want to apply, and then click **OK**.
  3. From the job history window, right-click the job whose log files you want to view, and then click **View Logs**.
  4. The contents of the log file related to the selected job history are displayed in the **Log File for Job n** window.
- 

[Back to Top](#)



# QR Disaster Recovery Solution for Building a UNIX Standby Oracle Server with Log Recovery

---

Overview

Configuration

- Prepare the Catalog Database Server
- Prepare the Production Server
- Prepare the Standby Server
- Configure the Quick Recovery Agent and Oracle iDataAgent

Bring the Standby Server online

Reinstall the Agent software on the new Production Server

Appendices

- Rename a Solaris Machine
- 

## OVERVIEW

In this setup, there are three servers, a *Production Server* which runs the Oracle Database, a *Catalog Server* with *Recovery Catalog Database* configured, and a *Standby Server*.

The Quick Recovery Agent takes a snapshot of the running instance. All tablespaces of the database are put in backup mode before taking the snapshot (Point-In-Time-Image) of the data partitions; as soon as the snapshot is taken, all the tablespaces are put back in open mode, then, after archiving the online redo logs, a snapshot of the archive log partition(s) is taken. This ensures that any transactions that occur during the backup mode of the tablespaces are archived. This archived redo log can be used for later Quick Recovery.

The *Production Server* contains the source (Primary) volumes, which are copied and updated incrementally via the LAN Copy Manager to QR Volumes on the *Standby Server*. The data is therefore physically ready to use (no lengthy restore from tape or other media is necessary) on the *Standby Server*. The *Standby Server* should have the same version of the operating system and the same version of Oracle installed on the same path as the *Production Server*. The *Catalog Server* hosts a Recovery Catalog Database used to recover the Production database in case of failure.

In cases where the *Production Server* suffers a failure, or requires downtime, the system administrator can remove the *Production Server* from the network and then add the *Standby Server* to the network (with the same name and configuration as the disconnected *Production Server*). After the *Standby Server* is up, a Recover using Archive Logs can be performed before opening the database. This will recover the data changed after the last successful QR Volume creation job and up to the last successful Archive Logs backup.

The key to a successful QR Recovery is properly configuring the *Production* and *Standby Servers*. This ensures that the database(s) can quickly and easily be run from the QR volumes on the *Standby Server*.

This procedure is written for bringing up one instance on the *Standby Server*. If multiple instances are present on one client we strongly recommend creating a separate subclient for each instance and following the same procedure for each instance individually.

It should be possible to accomplish this procedure between different versions of operating systems and Oracle databases, as well as different locations for Oracle installations. However, due to the significant complexity of dealing with all the additional factors involved, this document only covers the most common scenario.

---

## CONFIGURATION

The following sections discuss preparing the *Catalog Database*, *Production*, and *Standby Servers* as well as configuring the QR Agent and Oracle iDataAgent. The basic workflow is described below. For detailed instructions on installation and configuration options, see Quick Recovery Agent.

### PREPARE THE CATALOG DATABASE SERVER

To restore and recover the data that has changed after the last successful QR volume creation job, you must configure the Recovery Catalog. For detailed descriptions on how to create and configure the Recovery Catalog for the version of Oracle you are using, refer to Oracle Documentation.

- Create an Oracle Instance on the server you are planning to use as the *Recovery Catalog Server*.
- Create a schema and catalog using RMAN.
- Verify that the Target Database (*Production Server* database) is accessible using TNSPING. Edit TNSNAMES to correct any connection problems.
- Continue with the configuration of the *Production Server*. Later, you will need to register the Target Database with the Recovery Catalog from *Production Server*.

## PREPARE THE PRODUCTION SERVER

The Quick Recovery Agent supports Oracle 9.2 and higher versions. The database must be in "archive log mode" for QR Volume Creation. We recommend that data files, archive logs, control files, and online redo logs each reside on different partitions (different mount points), and that none of these four types of files reside on a partition with any of the others.

Following is a sample configuration of an Oracle installation on a Solaris 2.8 machine, which has an Oracle instance, "CV", running. Oracle requires that the *Standby* database be on the same OS level as that on the Production Server.

| Mount Point   | Description                                    |
|---------------|------------------------------------------------|
| /oracle901    | Contains the Oracle installed binaries         |
| /oracle901    | Is the ORACLE_HOME of instance CV              |
| /cv_data      | Contains all the data files of instance CV     |
| /cv_archlog   | Contains all the archived logs of instance CV  |
| /cv_control1  | Contains current control file 1 of instance CV |
| /cv_control02 | Contains current control file 2 of instance CV |
| /cv_redologs  | Contains online redo logs of instance CV       |

The redo logs and control files are not backed up by the QR Volume creation jobs since they are constantly changing. In order to take a consistent control file backup, Oracle provides specific commands to back up online control files. It is strongly recommended to have each database's data and archived logs on different volumes exclusively for each instance. (Example: /cv\_data or /cv\_archlog is on exclusive volume for instance CV. It should not have data or archive logs of any other instance running on that machine).

1. If the *Production Server* is not already installed with Oracle, install the desired version of Oracle and apply any required patches.
2. Create the required Oracle database.  
**NOTE:** When creating the instances or when installing Oracle, record your Oracle database configuration and storage locations so that Oracle can be installed identically on the *Standby Server*.
3. Register the database with the *Recovery Catalog*. For more information refer to the Oracle procedure for creating and using *Recovery Catalog*.
4. Install the QR Agent. If the installation detects that you have not already installed the Base software and the File System *iDataAgent* prior to installing the QR Agent, you will be prompted to install them.
5. Install the Oracle *iDataAgent* and configure the instance from the CommCell Console. Create a subclient that will back up the logs and control files. You will need to re-link the database by running the script `Ora_install.sh` (in the *iDataAgent* folder under the software installation) if you didn't specify the instance during the Oracle *iDataAgent* installation.  
**NOTE:** While configuring the instance, specify the connect string to the Recovery Catalog Database.
6. Shut down the Oracle instance to configure the disks used by Oracle as CXBF devices. Use Volume Explorer to perform this operation. If the devices are mounted you can select the **Unmount before Configure** option from the Volume Configuration dialog. The device will automatically be mounted back. Select Detect Volumes from Volume Explorer and verify all mount points used by the Oracle instance are configured and mounted.
7. Start up the Oracle instance on the *Production Server*.

## PREPARE THE STANDBY SERVER

Before installing any software on the *Standby Server*, consider the following:

- The *Standby Server* should have adequate resources to host the Oracle database(s).
- The *Standby Server* should have the same operating system version as the *Production Server*.
- The scratch volumes on the *Standby Server* must be of equal or greater size than the Primary volumes on the *Production Server*.

Install the same version of Oracle in the *Standby Server* as is installed on the *Production Server*. You must use the same Oracle user ID, Oracle group ID, and the same installation path of Oracle as the *Production Server*.

1. Set up the Oracle instance with the same instance name and configuration as that of the Production Server:
  - a. Create the required directories (`create`, `cdump`, `bdump`, `udump`, and `pfile`) for the instance in the *Standby Server* matching the *Production Server*. (Example location: `$ORACLE_BASE/admin/<ORACLE_SID>/cdump`)
  - b. Oracle database may be using `spfile` or `init` file for storing the initialization parameters. Copy the `init/spfile` (`init<ORACLE_SID>.ora` or `spfile<ORACLE_SID>`) from the *Production Server* to the *Standby Server*. Create the necessary links for the `init` file if required. (Example location of these files: `$ORACLE_HOME/dbs`)
  - c. Copy the password file (example: `orapw<ORACLE_SID>`) of the instance from the *Production Server* to the *Standby Server*. (Example location of this file: `$ORACLE_HOME/dbs`)
  - d. Copy the `TNSNAMES` file from the *Production Server*. It contains information about the Recovery Catalog Database. Copy any other files specific to your configuration that you may need to connect to Recovery Catalog. Use `TNSPING` to verify the connection.  
**NOTE:** Run the Oracle `netmgr` tool (from the machine console) if Oracle has a problem resolving the instance name.
  - e. On the *Production Server*, note the name and location of each control file. This step is required to copy back the backup control file to each control file

location when bringing up the *Standby Server*.

Example: `/cv_control1/control101.ct1, /cv_control2/control102.ct1, /cv_control3/control103.ct1`

**NOTE:**

It is very important to follow Step 1 precisely. These files are used for bringing up the database on the *Standby Server* in the event the *Production Server* goes down. Failure to properly configure the *Standby Server* will result in not being able to bring up the *Standby Server* successfully.

2. Install the MediaAgent and Quick Recovery Agent software on the *Standby Server*. If the installation detects that you have not already installed the Base software and the File System *iDataAgent* prior to installing the QR Agent, you will be prompted to install them.
3. Start up the instance using the **nomount** option.
4. Install the Oracle *iDataAgent* and configure the instance from the CommCell Console. Specify the connect string to catalog database. You may need to re-link the instance as mentioned above.
5. Shut down the Oracle instance on the *Standby Server*.
6. Use Volume Explorer to configure all volumes you are planning to use in the Volume Scratch Pool on the *Standby Server* as CXBF devices. Detect volumes when finished.

---

## CONFIGURE THE QUICK RECOVERY AGENT AND ORACLE /DATAAGENT

1. Configure the scratch volume pool. The scratch volume pool for this recovery scenario should consist only of the disk resources attached to the *Standby Server*. The QR volumes and incremental updates will be written to these volumes. The on-line help contains detailed steps for this operation.
2. Create a QR Policy. Select the LAN Copy Manager (LANVolCopy) on the *Standby Server* as the copy manager, and associate this QR Policy with the scratch volume pool that was created in the previous step. Multiple QR Policies may be associated with a given scratch volume pool.
3. In the **Authentication** tab of the QR Agent Properties dialog box, add the selected Oracle instance.
4. Create the QR Agent subclient(s) on the *Production Server*. Please note the following:
  - o The QR Agent manages Oracle databases at the instance level. Therefore, different databases can be distributed across multiple subclients.
  - o Add the Oracle instance to a QR subclient through the **Add App** option located in the *Subclient Contents* property box. Verify all Oracle volumes containing data files and archive logs are listed. Disks containing redo log files and control files will not be listed unless they reside on the same volume as data or archive logs.
5. Create an Oracle *iDataAgent* subclient on the *Production Server*. As subclient content, select only Archive Logs. Deselect data files.
6. Start or schedule the QR Volume creation job. In this particular scenario, incremental updates should be scheduled. Also schedule backups for Archive Logs using the Oracle *iDataAgent*. In general, Archive Log backups should occur often, in between QR Volume creation jobs. There should be one more Archive Log backup after the last QR Volume creation job. These will be used to recover the database to a point-in-time, which is after the last snap.

From the *QR Volume Creation Advanced Options* dialog box, assign each volume on the *Production Server* to its corresponding destination volume on the *Standby Server*. The mount path of each *standby* volume should be the same as its counterpart *production* volume. For each source raw partition on the *Production Server*, the corresponding destination raw partition on the *Standby Server* should be selected in the *QR Volume Creation Advanced Options* dialog box.

If any raw partitions are used for data files or control files, you must create the necessary links for appropriate devices on the *Standby Server* before QR Volume Creation. The soft link pointing to the *destination* raw partition should be the same as that of the *source* raw partition.

For example, if the production database has a data file linked to a raw volume (data file) as follows:

```
/ora_data/<ORACLE_SID>/raw.dbf -> /dev/cxbf/rdisk/ct1d1s1 (source)
```

then the links to the destination volume should be created on the *Standby Server* as follows (after selecting the destination volume in the *Advanced* tab and before starting the QR Volume creation process):

```
/ora_data/<ORACLE_SID>/raw.dbf -> /dev/cxbf/rdisk/c2t1d1s1 (destination)
```

7. To test this configuration, allow some scheduled updates to complete. If necessary, execute a few test Oracle transactions between updates to verify that changed/new blocks are being copied to the QR volumes.

---

## BRING THE STANDBY SERVER ONLINE

In cases where the *Production Server* suffers a failure, or requires downtime, the *Standby Server* can quickly be brought on-line to host the database from the QR Volumes you have created. The following steps must be taken to add the *Standby Server* to the network.

1. Shut down the *Production Server* and/or remove it completely from the network.
2. Mount all Oracle volumes on the *Standby Server*. Set ownership to Oracle user if needed.
3. There is a `backup.ct1.galaxy` (Backup Controlfile) file located in one of the archive log destinations of the instance on the *Standby Server* (Example location: `/ora_logs/admin/sid/arch`)

Copy the `backup.ct1.galaxy` file to all the control file location as noted in Step 1 of "Prepare the Standby Server".

Example:

```
cp backup.ct1.galaxy /cv_control1/control01.ct1
cp backup.ct1.galaxy /cv_control2/control02.ct1
cp backup.ct1.galaxy /cv_control3/control03.ct1
```

If the control file is on a raw device, then the above example would change to:

```
dd if=backup.ct1.galaxy of=/cv_control1/controlraw.ct1
```

#### NOTES:

- o This is the copy of control files backed up by the QR Agent.
  - o Oracle user should copy these control files to the appropriate location.
  - o Before proceeding to Step 4, set up the Oracle environment (such as `ORACLE_HOME` and `ORACLE_SID`) on the destination machine to match that of the *Production Server*.  
(Example: `export ORACLE_HOME=/oracle901` and `ORACLE_SID=CV`)
4. Start the database in the **mount** mode after connecting as the `sysdba` user (Example: `sys/password as sysdba`):

```
SQL> startup mount
```

5. Restore Archive logs:
- a. Using the CommCell Console, select the Oracle *iDataAgent* on the *Production Server* and select **Browse**.
  - b. Select the Oracle instance and click **Recover all selected**.
  - c. Deselect **Restore Data** and select **Restore Archive logs**. Click **Advanced Options** and select **Restore to End** (clear **Start**). Select a time before or equal to the last Archive logs backup and click **OK**.
  - d. Change the restore location to the *Standby Server*. Provide **Recovery Catalog** login information and click **OK** to start the restore.
6. Recover the Database:
- a. Using the CommCell Console select the Oracle *iDataAgent* on the *Production Server* and select **Browse**.
  - b. Select the Oracle instance and click **Recover all selected**.
  - c. Deselect **Restore Data** and select **Recover**. Click **Advanced Options** and select **Recover to End** (clear **Start**). Select a time before or equal to the last Archive logs you restored. Alternately, you can use SCN.
  - d. Change the recover location to the *Standby Server*. Provide **Recovery Catalog** login information and click **OK** to start the recovery.
  - e. At the end of the recovery, the database will be altered to OPEN. Verify in the RMAN Log that the operation was successful.
7. Rename the *Standby Server* to match the name of the *Production Server* that has been removed from the network. (Please refer to Solaris user documentation for renaming the machine. One possible way is to use the **sys-unconfig** command, then delete DNS entries; when the machine restarts, use the name and IP of the *Production Server*). If necessary, change the IP Address to match the one on the *Production Server*. Using the old IP with a new name could cause problems with name resolution. Add the *Standby Server* to the network if necessary, and configure TCP/IP and any other applicable network settings.

After bringing up the *Standby Server*, verify that the destination volumes are mounted back. These are the volumes that were present in the scratch volume pool while creating QR volumes. Also verify Oracle user is an owner of all mount points, directories and files. If necessary, change the ownership to the appropriate Oracle user and group; otherwise your instance might fail to access the files.

#### NOTE:

Do not try to detect volumes in Volume Explorer on either the *Production* or the *Standby Server*. The Oracle recovery procedure is not accomplished with Volume Explorer.

#### OPTIONAL:

Add the destination volumes and their mount points to `/etc/vfstab` on the *Standby Server*, to ensure the volumes will be mounted back in the event of a re-boot.

After the successful completion of Step 7, the database should be in OPEN mode and ready to use.

#### NOTE:

- If you receive any messages that files cannot be accessed or access is denied during the startup process, check again that you have mounted all volumes, copied the files, and set the ownership to Oracle user.

## REINSTALL THE AGENT SOFTWARE ON THE NEW PRODUCTION SERVER

Once your Oracle database is up and running and the computer name changed to the Production Server name, you must reinstall the Agent software if you plan

to perform backups from this machine. This is because all configuration settings still refer to the old name of the machine (*Standby Server Name*).

1. Uninstall all Agent software from the new Production Server. Do not delete remaining client entries from the CommCell Console.
2. Reinstall the same modules you had on the *Production Server*. They will be installed with the new machine name (*Production Server one*)
3. If you have the new *Standby Server* ready, configure it according to the section above, "Prepare the Standby Server."
4. Proceed with the section above, "Configure the QR Agent and Oracle iDataAgent."

**NOTE:**

- In most scenarios, the machine that used to be *Production Server* could be easily turned into a *Standby Server* after its maintenance or repair is completed; follow the procedures above.
- 

## APPENDICES

---

### RENAME A SOLARIS MACHINE

There is a utility called **sys-unconfig** which is used for resetting the network configuration. After using it, reboot the machine and enter the new IP, hostname, etc. You can achieve the same effect by editing the listed files below and then rebooting.

`/etc/hosts`

This is where you specify the IP address of your hostname. Change the hostname here to match your host's new name. It should be the same as what you specify in `/etc/nodename`.

`/etc/nodename`

This is just like `/etc/HOSTNAME` in Linux. It defines the real name of the host. Simply change it to whatever you want to call your host.

`/etc/hostname.hme0` (or other interface name)

Change these to line up with what you specified in `/etc/hosts`.

`/etc/net/tic*/hosts`

Change everything in here to line up with the files above.

`/etc/resolv.conf`

Just like Linux. This is where you specify your DNS servers and domain resolution information.

`/etc/defaultrouter`

Enter the IP address of the default router for your Solaris host. Note that Sun does not create this file by default, nor is there any other location where you may specify a default route.

[Back to Top](#)

---

# QR Disaster Recovery Solution for Building a UNIX Standby Oracle Server

Overview

Configuration

- Prepare the Production Server
- Prepare the Standby Server
- Configure the Quick Recovery Agent

Bring the Standby Server online

Appendices

- Rename a Solaris Machine

## OVERVIEW

In this scenario, the user has at least two servers, a *Production Server*, which runs the Oracle Database and a *Standby Server*.

The Quick Recovery Agent takes a snapshot of the running instance. All tablespaces of the database are put in backup mode before taking the snapshot (Point-In-Time-Image) of the data partitions; as soon as the snapshot is taken, all the tablespaces are put back in open mode, then, after archiving the online redo logs, a snapshot of the archive log partition(s) is taken. This ensures that any transactions happening during the backup mode of the tablespaces are archived. This archived redo log can be used for later Quick Recovery.

The *Production Server* contains the source (Primary) volumes, which are copied and updated incrementally via the LAN Copy Manager to QR Volumes on the *Standby Server*. The data is therefore physically ready to use (no lengthy recovery from tape or other media is necessary) on the *Standby Server*. The *Standby Server* should have the same version of the Operating system and the same version of Oracle installed on the same path as the *Production Server*.

In cases where the *Production Server* suffers a failure, or requires downtime, the system administrator can remove the *Production Server* from the network and then add the *Standby Server* to the network (with the same name and configuration as the disconnected *Production Server*).

The key to a successful *QR Agent* recovery is properly configuring the *Production* and *Standby Servers*. This ensures that the database(s) can quickly and easily be run from the QR volumes on the *Standby Server*.

This procedure is written for bringing up one instance on the *Standby Server*. If multiple instances are present on one client we strongly recommend creating a separate subclient for each instance and following the same procedure for each instance individually. This procedure has been tested on Oracle 9.2 database with Sun Solaris 8 and the Sparc platform. The procedure for other versions of Oracle and the OS may vary slightly.

## CONFIGURATION

The following sections discuss preparing the *Production* and *Standby Servers* as well as configuring the QR Agent. The basic workflow is described below. For detailed instructions on installation and configuration options, see Quick Recovery Agent.

### PREPARE THE PRODUCTION SERVER

The Quick Recovery Agent supports Oracle 9.2. The database must be in "archive log mode" for QR Volume Creation. We recommend that data files, archive logs, control files, and online redo logs each reside on different partitions (different mount points), and that none of these four types of files reside on a partition with any of the others.

Following is a sample configuration of an Oracle 9.2 installation on a Solaris 2.8 machine, which has an Oracle instance, "CV", running. Oracle requires that the *Standby* database be on the same OS level as that on the *Production Server*.

| Mount Point   | Description                                    |
|---------------|------------------------------------------------|
| /oracle901    | Contains the Oracle 9.2 installed binaries     |
| /oracle901    | Is the ORACLE_HOME of instance CV              |
| /cv_data      | Contains all the data files of instance CV     |
| /cv_archlog   | Contains all the archived logs of instance CV  |
| /cv_control1  | Contains current control file 1 of instance CV |
| /cv_control02 | Contains current control file 2 of instance CV |
| /cv_redologs  | Contains online redo logs of instance CV       |

The redo logs and control files are not backed up since they are constantly changing. In order to take a consistent control file backup, Oracle provides specific commands to back up online control files. It is strongly recommended to have each database's data and archived logs on different volumes exclusively for each instance. (Example: /cv\_data or /cv\_archlog is exclusively for instance cv. It should not have data or archive logs of any other instance running on that

machine).

1. If the *Production Server* is not already installed with Oracle 9.2, install the desired version of Oracle and apply any required patches.
2. Create the required Oracle database.  
**NOTE:** When creating the instances or when installing Oracle, record your Oracle database configuration and storage locations so that Oracle can be installed identically on the *Standby Server*.
3. Install the QR Agent. If the installation detects that you have not already installed the Base software and the File System *iDataAgent* prior to installing the QR Agent, you will be prompted to install them.
4. Shut down the Oracle instance to configure the disks used by Oracle as CXBF devices. Use *Volume Explorer* to perform this operation. If the devices are mounted, you can select the **Unmount Before Configure** option from the *Volume Configuration* dialog. The device will be mounted back automatically. Select **Detect Volumes** from *Volume Explorer* and verify that all mount points used by the Oracle instance are configured and mounted.
5. Start up the Oracle instance on the *Production Server*.

---

## PREPARE THE STANDBY SERVER

Before installing any software on the *Standby Server*, consider the following:

- The *Standby Server* should have adequate resources to host the Oracle database(s).
- The *Standby Server* should have the same operating system version as the *Production Server*.
- The *scratch volumes* created on the *Standby Server* must be of a size equal to or greater size than the *Primary* volumes on the *Production Server*. It is recommended that the *scratch volumes* be larger than the *Primary* volumes.

Install the same version of Oracle as is installed on the *Production Server*. You must use the same Oracle user ID, Oracle group ID, and the same installation path of Oracle as the *Production Server*.

1. Set up the Oracle instance with the same instance name and configuration as that of the *Production Server*:
  - a. Create the required directories (`admin`, `cdump`, `bdump`, `udump`, and `pfile`) for the instance in the *Standby Server* matching the *Production Server*.  
Example location: `$ORACLE_HOME/admin/<ORACLE_SID>/cdump`
  - b. Oracle database may be using `spfile` or `init` file for storing the initialization parameters. Copy the `init/spfile` (`init<ORACLE_SID>.ora` or `spfile<ORACLE_SID>`) from the *Production Server* to the *Standby Server*. Create the necessary links for the `init` file if required.  
Example location of these files: `$ORACLE_HOME/dbs`
  - c. Copy the password file (example: `orapw<ORACLE_SID>`) of the instance from the *Production Server* to the *Standby Server*.  
Example location of this file: `$ORACLE_HOME/dbs`
  - d. On the *Production Server*, note the name and location of each control file. This step is required to copy back the backup control file to each control file location when bringing up the *Standby Server*.  
Example: `/cv_control1/control01.ct1, /cv_control2/control02.ct1, /cv_control3/control03.ct1`

### NOTE:

It is very important to follow Step 1 precisely. These files are used for bringing up the database on the *Standby Server* in the event the *Production Server* goes down. Failure to properly configure the *Standby Server* will result in not being able to bring up the *Standby Server* successfully.

2. Install the MediaAgent and Quick Recovery Agent software on the *Standby Server*. If the installation detects that you have not already installed the Base software and the File System *iDataAgent* prior to installing the QR Agent, you will be prompted to install them.
3. Use *Volume Explorer* to configure all volumes you are planning to use in the *Volume Scratch Pool* on the *Standby Server*. Detect volumes when finished.

---

## CONFIGURE THE QUICK RECOVERY AGENT

1. Configure the scratch volume pool. The scratch volume pool for this recovery scenario should consist only of the disk resources attached to the *Standby Server*. The QR volumes and incremental updates will be written to these volumes. The on-line help contains detailed steps for this operation.
2. Create a QR Policy. Select the LAN Copy Manager (LANVolCopy) on the *Standby Server* as the copy manager, and associate this QR Policy with the scratch volume pool that was created in the previous step. Multiple QR Policies may be associated with a given scratch volume pool.
3. In the **Authentication** tab of the QR Agent Properties dialog box, add the selected Oracle instance.
4. Create the QR Agent subclient(s) on the *Production Server*. Please note the following:
  - The QR Agent manages Oracle databases on the instance level. Therefore, different databases can be distributed across multiple subclients.
  - Add the Oracle instance to a QR subclient through the **Add App** option located in the *Subclient Contents* property box. Verify all Oracle volumes containing data files and archive logs are listed. Disks containing redo log files and control files will not be listed unless they reside on the same volume as data or archive logs
5. Start or schedule the QR Volume creation job. In this particular scenario, incremental updates should be scheduled. Please note that we do not support archive logs backup with the *Standby Server*. The QR incremental updates already take care of archive logs that were created after the last QR update.

From the *QR Volume Creation Advanced Options* dialog box, assign each volume on the *Production Server* to its corresponding destination volume on the

*Standby Server*. The mount path of each *standby* volume should be same as its counterpart *production* volume. For each source raw partition on the *Production Server*, the corresponding destination raw partition on the *Standby Server* should be selected in the *QR Volume Creation Advanced Options* dialog box.

If any raw partitions are used for data files or control files, you must create the necessary links for appropriate devices on the *Standby Server* before QR Volume creation. The *soft link* pointing to the destination raw partition should be the same as that of the source *raw* partition.

For example, if the production database has a data file linked to a raw volume (data file) as follows:

```
/ora_data/<ORACLE_SID>/raw.dbf -> /dev/cvbf/rdisk/ct1t1d1s1 (source)
```

then the links to the destination volume should be created on the *Standby Server* as follows (after selecting the destination volume in the *Advanced* tab and before starting the QR Volume creation process):

```
/ora_data/<ORACLE_SID>/raw.dbf -> /dev/cvbf/rdisk/c2t1d1s1 (destination)
```

6. To test this configuration, allow some scheduled updates to complete. If necessary, execute a few test Oracle transactions between updates to verify that changed/new blocks are being copied to the QR volumes.

## BRING THE STANDBY SERVER ONLINE

In cases where the *Production Server* suffers a failure, or requires downtime, the *Standby Server* can quickly be brought on-line to host the database from the QR Volumes you have created. The following steps must be taken to add the *Standby Server* into the network.

1. Shut down the *Production Server* and/or remove it completely from the network.
2. Rename the *Standby Server* to match the name of the *Production Server* that has been removed from the network. (Please refer to Solaris user documentation for renaming the machine). If necessary, change the IP Address to match the one on the *Production Server*. Using the old IP with a new name could cause problems with name resolution. Add the *Standby Server* to the network if necessary, and configure TCP/IP and any other applicable network settings.

After bringing up the *Standby Server*, verify that the destination volumes are mounted back. These are the volumes that were present in the scratch volume pool while creating QR volumes. Also verify Oracle user is an owner of all mount points, directories and files. If necessary change the ownership to the appropriate Oracle user and group; otherwise your instance might fail to access the files.

### NOTE:

Do not try to detect volumes in Volume Explorer on either the *Production* or the *Standby Server*. The Oracle recovery procedure is manual, and is not accomplished with *Volume Explorer*, the *QR Agent*, or the *Oracle iDataAgent*.

### OPTIONAL:

Add the destination volumes and their mount points to */etc/vfstab* on the *Standby Server*, to ensure the volumes will be mounted back in the event of a re-boot.

3. There is a *backup.ct1.galaxy* (Backup Controlfile) file located in one of the archive log destinations of the instance on the *Standby Server*. (Example location: */ora\_logs/admin/sid/arch*)

Copy the *backup.ct1.galaxy* file to all the control file locations as noted down in Step 1 of "Prepare the Standby Server".

Example:

```
cp backup.ct1.galaxy /cv_control1/control01.ct1
```

```
cp backup.ct1.galaxy /cv_control2/control02.ct1
```

```
cp backup.ct1.galaxy /cv_control3/control03.ct1
```

If the control file is on a raw device, then the above example would change to:

```
dd if=backup.ct1.galaxy of=/cv_control1/controlraw.ct1
```

### NOTES:

- o This is the copy of control files backed up by QR Agent.
- o Oracle user should copy these control files to the appropriate location.
- o Before proceeding to Step 4, set up the Oracle environment (such as *ORACLE\_HOME* and *ORACLE\_SID*) on the destination machine to match that of the *Production Server*.  
(Example: *export ORACLE\_HOME=/oracle901 and ORACLE\_SID=CV*)

4. Start the database in the **mount** mode after connecting as the *sysdba* user (Example: *sys/password as sysdba*):

```
SQL> startup mount
```

5. Recover the database as follows:

```
o SQL> set autorecovery on
```

```
o SQL> recover database until cancel using backup controlfile;
```



```
O SQL> alter database open resetlogs;
```

After the successful completion of Step 5, the database is in OPEN mode and ready to use.

**NOTE:**

If you receive any messages that files cannot be accessed or access is denied during the startup process, check again that you have mounted all volumes, copied the files and set the ownership to Oracle user.

---

## APPENDICES

---

### RENAME A SOLARIS MACHINE

There is a utility called **sys-unconfig** which is used for resetting the network configuration. After using it, reboot the machine and enter the new IP, hostname, etc. You can achieve the same effect by editing the listed files below and then rebooting.

`/etc/hosts`

This is where you specify the IP address of your hostname. Change the hostname here to match your host's new name. It should be the same as what you specify in `/etc/nodename`.

`/etc/nodename`

This is just like `/etc/HOSTNAME` in Linux. It defines the real name of the host. Simply change it to whatever you want to call your host.

`/etc/hostname.hme0` (or other interface name)

Change these to line up with what you specified in `/etc/hosts`.

`/etc/net/tic*/hosts`

Change everything in here to line up with the files above.

`/etc/resolv.conf`

Just like Linux. This is where you specify your DNS servers and domain resolution information.

`/etc/defaultrouter`

Enter the IP address of the default router for your Solaris host. Note that Sun does not create this file by default, nor is there any other location where you may specify a default route.

[Back to Top](#)

---

# QR Disaster Recovery Solution for Building a Windows Standby Oracle Server

## Overview

### Configuration

- Prepare the Production Server
- Prepare the Standby Server
- Configure the Quick Recovery Agent

Bring the Standby Server online

## OVERVIEW

This document describes the procedure necessary to create a *Standby Exchange Server* in the event that a *Production Exchange Server* is temporarily or permanently damaged. These procedures are for an environment with at least two servers, a *Production Server*, which runs the Oracle Database, and a *Standby Server*.

*QR Agent* takes a snapshot of the running instance. All tablespaces of the database are put in backup mode before taking the snapshot (*Point-In-Time-Image*) of the data partitions; as soon as the snapshot is taken, all the tablespaces are put back in open mode, then, after archiving the online redo logs, a snapshot of the archive log partition(s) is taken. This ensures that any transactions happening during the backup mode of the tablespaces are archived. This archived redo logs can be used for later Quick Recovery.

The *Production Server* contains the source (Primary) volumes, which are copied and updated incrementally via the LAN Copy Manager to QR Volumes on the *Standby Server*. The data is therefore physically ready to use (no lengthy recovery from tape or other media is necessary) on the *Standby Server*. The *Standby Server* should have the same version of the Operating system and the same version of Oracle installed on the same path as in the *Production Server*.

In cases where the *Production Server* suffers a failure, or requires downtime, the system administrator can remove the *Production Server* from the network and then add the *Standby Server* to the network (with the same name and configuration as the disconnected *Production Server*).

The key to a successful QR Agent recovery is properly configuring the *Production* and *Standby Servers*. This ensures that the database(s) can quickly and easily be run from the QR volumes on the *Standby Server*.

This procedure is written for bringing up one instance on the *Standby Server*. If multiple instances are present on one client we strongly recommend creating a separate subclient for each instance and following the same procedure for each instance individually.

## CONFIGURATION

The following sections discuss preparing the *Production* and *Standby Servers* as well as configuring the QR Agent. The basic workflow is described below. For detailed instructions on installation and configuration options, see Quick Recovery Agent.

### PREPARE THE PRODUCTION SERVER

The Quick Recovery Agent supports Oracle 9.2 and higher versions. The database must be in "archive log mode" for QR Volume Creation. We recommend that data files, archive logs, control files, and online redo logs each reside on different partitions (different mount points), and that none of these four types of files reside on a partition with any of the others.

Following is a sample configuration of an Oracle installation on a Windows 2000 machine, which has running an Oracle instance, "CV". Oracle requires that the *Standby* database to be on the same OS level as that on the *Production Server*.

| Mount Point    | Description                                   |
|----------------|-----------------------------------------------|
| V:\oracle92    | Contains the Oracle installed binaries        |
| V:\oracle92    | Is the ORACLE_HOME of instance CV             |
| W:\cv_data     | Contains all the data files of instance CV    |
| X:\cv_archlog  | Contains all the archived logs of instance CV |
| Y:\cv_control  | Contains current control file of instance CV  |
| Z:\cv_redologs | Contains online redo logs of instance CV      |

The redo logs and control files are not backed up since they are constantly changing. In order to take a consistent control file backup, Oracle provides specific commands to back up online control files. It is strongly recommended to have each database's data and archived logs on different volumes exclusively for each instance. (Example: W:\cv\_data or X:\cv\_archlog is exclusively for instance CV. It should not have data or archive logs of any other instance running on that machine).

1. If the *Production Server* is not already installed with Oracle, install the desired version of Oracle and apply any required patches.

2. Create the required Oracle database.

**NOTE:** When creating the instances or when installing Oracle, record your Oracle database configuration and storage locations so that Oracle can be installed identically on the Standby Server.

3. Install the QR Agent, Oracle *iDataAgent*, and MediaAgent. If the installation detects that you have not already installed the *File System iDataAgent* prior to installing the *QR Agent*, you will be prompted to install it.

---

## PREPARE THE STANDBY SERVER

Before installing any software on the *Standby Server*, consider the following:

- The *Standby Server* should have adequate resources to host the Oracle database(s).
- The *Standby Server* should have the same operating system version as the *Production Server*.
- The *scratch volumes* created on the *Standby Server* must be of equal or greater size than the Primary volumes on the *Production Server*. It is recommended that the *scratch volumes* be larger than the Primary volumes.

Install the same version of Oracle as is installed on the *Production Server*. You must have the same Oracle user ID, Oracle group ID, and the same installation path of Oracle as the *Production Server*.

1. Set up the Oracle instance with the same instance name and configuration as that of the *Production Server*:
  - a. Verify the required directories (*admin*, *cdump*, *bdump*, *udump*, and *pfile*) for the instance in the *Standby Server* match the *Production Server*. (Example location: %Install Dir%\admin%\ORACLE\_SID\cdump)
  - b. Oracle database may be using *spfile* or *init* file for storing the initialization parameters. Copy the *init\spfile* (*init*<ORACLE\_SID>.ora or *spfile*<ORACLE\_SID>) from the *Production Server* to the *Standby Server*. (Example location of these files: %Install Dir%\database)
  - c. Copy the password file (example: *orapw*<ORACLE\_SID>) of the instance from the *Production Server* to the *Standby Server*. (Example location of this file: %Install Dir%\database)
  - d. On the *Production Server*, note the name and location of each control file. This step is required to copy back the backup control file to each control file location when bringing up the *Standby Server*.

Example:

```
%Install Dir%\cv_control\control01.ctl
%Install Dir%\cv_control\control02.ctl
%Install Dir%\cv_control\control03.ctl
```

**NOTE:** It is very important to follow Step 1 precisely. These files are used for bringing up the database on the *Standby Server* in the event the *Production Server* goes down. Failure to properly configure the *Standby Server* will result in not being able to bring up the *Standby Server* successfully.

2. Install the MediaAgent and Quick Recovery Agent software on the *Standby Server*. If the installation detects that you have not already installed the *File System iDataAgent* prior to installing the *QR Agent*, you will be prompted to install it.
3. Use *Volume Explorer* to configure all volumes you are planning to use in the *Volume Scratch Pool* on the *Standby Server*.

---

## CONFIGURE THE QUICK RECOVERY AGENT

1. Configure the scratch volume pool. The scratch volume pool for this recovery scenario should consist only of the disk resources attached to the *Standby Server*. These volumes must have the same mount point as the source volumes on the *Production Server*. The QR volumes and incremental updates will be written to these volumes.
2. Create a QR Policy. Select the LAN Copy Manager (LANVolCopy) on the *Standby Server* as the copy manager, and associate this QR Policy with the scratch volume pool that was created in the previous step. Multiple QR Policies may be associated with a given scratch volume pool.
3. Select the **Authentication** tab in the QR Agent Properties dialog box and add the Oracle instance.
4. Create the QR Agent subclient(s) on the *Production Server*. Please note the following:
  - The QR Agent manages Oracle databases on the instance level. Therefore, different databases can be distributed across multiple subclients.
  - Add the Oracle instance to a QR subclient through the *Add App* option located in the *Subclient Contents* property box. Verify all Oracle volumes containing data files and archive logs are listed. Disks containing redo log files and control files will not be listed unless they reside on the same volume as data or archive logs.
5. Schedule the QR Volume creation job. In this particular scenario, incremental updates should be selected. Please note that we do not support archive logs backup with the *Standby Server*. The QR incremental updates already take care of archive logs that were created after the last QR update.

From the *QR Volume Creation Advanced Options* dialog box, assign each volume on the *Production Server* to its corresponding destination volume on the *Standby Server*. The mount path of each standby volume must be same as its counterpart production volume.

---

## BRING THE STANDBY SERVER ONLINE

In cases where the *Production Server* suffers a failure, or requires downtime, the *Standby Server* can quickly be brought on-line to host the database from the QR Volumes you have created. The following steps must be taken to add the *Standby Server* into the network.

1. Shut down the *Production Server* and/or remove it completely from the network.
2. Rename the *Standby Server* to match the name of the *Production Server* that has been removed from the network. (Please refer to Windows user documentation for renaming the machine.) If necessary, change the IP Address to match the one on the *Production Server*. Add the *Standby Server* to the network, and configure any other applicable network settings.

After bringing up the *Standby Server*, verify that the destination volumes are mounted. These are the volumes that were present in the scratch volume pool while creating QR volumes. Also verify Oracle user is an owner of all mount points, directories and files. If necessary change the ownership to the appropriate Oracle user and group; otherwise your instance might fail to access the files.

**NOTE:** Do not try to detect volumes in Volume Explorer on either the *Production* or the *Standby Server* at this time. The Oracle recovery procedure is manual, and is not accomplished with *Volume Explorer*, the *QR Agent*, or the *Oracle iDataAgent*.

3. There is a `backup.ct1.galaxy` (Backup Controlfile) file located in one of the archive log destinations of the instance on the *Standby Server*. (Example location: `\ora_logs\admin\sid\arch`)

Copy the `backup.ct1.galaxy` file to all the control file location as in "Prepare the Standby Server".

Example:

```
copy backup.ct1.galaxy Y:\cv_control\control01.ct1
```

```
copy backup.ct1.galaxy Y:\cv_control\control02.ct1
```

```
copy backup.ct1.galaxy Y:\cv_control\control03.ct1
```

**NOTE:** Before proceeding to Step 4, set up the Oracle environment (such as `ORACLE_HOME` and `ORACLE_SID`) on the destination machine to match that of the *Production Server*.

(Example: `export ORACLE_HOME=/oracle901 and ORACLE_SID=CV`)

4. Start the database in the mount mode after connecting as the `sysdba` user (Example: `sys/password as sysdba`):

```
SQL> startup mount
```

5. Recover the database as follows:

```
SQL> set autorecovery on
```

```
SQL> recover database until cancel using backup controlfile;
```

```
SQL> alter database open resetlogs;
```

After the successful completion of Step 5, the database is in OPEN mode and ready to use.

#### NOTES:

If you receive any messages that files cannot be accessed or access is denied during the startup process, check again that you have mounted all volumes, copied the files and set the ownership to Oracle user.

[Back to Top](#)

---

# QR Disaster Recovery Solution for Building a Windows Standby SQL Server

---

Overview

Configuration

- Prepare the Production Server
- Prepare the Standby Server
- Configure the Quick Recovery Agent

Bring the Standby Server online

---

## OVERVIEW

This document describes the procedure to create a *Standby SQL Server* in the event that a *Production SQL Server* is temporarily or permanently unavailable.

These procedures are supported for SQL Server, on either a Windows 2000 or Windows 2003 machine. It has been certified with the Quick Recovery Agent and QSnap, using the QSnap snap engine for Windows 2000, and the VSS snap engine for Windows 2003. A familiarity with the functionality and configuration of the Quick Recovery Agent is necessary in order to properly conduct this procedure.

In the configuration stage, you will prepare the *Production* and *Standby Servers* for the procedure. The Quick Recovery Agent will copy the data from the *Production Server* to the *Standby Server*. To use the *Standby Server*, you will attach the database(s) that have been protected. SQL functionality will then continue using the *Standby Server*.

The document contains four sections. First, the *Production Server* is configured. Second, the *Standby Server* is configured. Third, the Quick Recovery Volumes that will be used on the *Standby Server* are created. Fourth, the application is brought online.

## ADVANTAGES

- Very fast recovery time. Recovery time is limited to the time taken to attach the SQL databases.
  - Recovery time does not include a lengthy restore from tape or other media.
  - There is no need to restore the system state, reinstall SQL, or its Service Packs.
- 

## CONFIGURATION

The following sections discuss preparing the *Production* and *Standby Servers* as well as configuring the QR Agent. The basic workflow is described below. For detailed instructions on installation and configuration options, see Quick Recovery Agent.

### PREPARE THE PRODUCTION SERVER

This procedure assumes that the *Production Server* has already been installed with SQL Server with the latest service packs or patches that may be needed.

1. Record your SQL configuration and storage locations so that SQL can be installed identically on the *Standby Server*.
2. Install the Quick Recovery Agent and QSnap on the *Production Server*. If you want to use VSS as your snap engine, install the VSS Enabler in addition to the other products already mentioned. If necessary, install the CommCell Console as a Stand-Alone Application.

### PREPARE THE STANDBY SERVER

Before installing any software on the *Standby Server*, verify the following:

- The *Standby Server* has adequate resources to host the SQL database(s).
  - The QR volumes that will be created on the *Standby Server* must be equal to or greater than the size of the Primary volumes on the *Production Server*.
1. Install SQL Server and apply any service packs or patches. The location of the install should be the same as the *Production Server*.
  2. Set up the SQL Server with the same instance names and configuration as the *Production Server*.
  3. Install the Quick Recovery Agent, QSnap, and MediaAgent on the *Standby Server*. If necessary, install CommCell Console as a Stand-Alone Application.

### CONFIGURE THE QUICK RECOVERY AGENT

1. Configure the scratch volume pool. The scratch volume pool for this recovery scenario should consist of the disk resources attached to the *Standby Server*. The QR volumes and incremental updates will be written to these volumes. For more information, see Scratch Volume Pools.

2. Create a QR Policy. Select the LAN Copy Manager (LANVolCopy) on the *Standby Server* as the copy manager, and associate this QR Policy with the scratch volume pool that was created in the previous step.
3. Create the QR Agent subclient(s) on the *Production Server*. Please note the following:
  - o The QR Agent manages SQL databases on the database level. Therefore, all the volumes associated with a given database (*mdf*, *ndf*, *ldf*) should be selected as subclient content.
  - o Adding volumes that contain SQL data to a QR subclient is accomplished through the *Add App* button located in the *Subclient Content* property box.
4. Start or schedule the QR incremental update job.
  - a. From the *QR Volume Creation Advanced Options* dialog box, assign each volume on the *Production Server* to its corresponding destination volume on the *Standby Server*.
  - b. Set the mount path of each production volume to the drive letter of the corresponding volume on the *Standby Server*.

The frequency of the incremental updates determines how often new and changed blocks are copied to the QR volumes. The database will be recovered to its state at the time of the last incremental update.

---

## BRING THE STANDBY SERVER ONLINE

In cases where the *Production Server* suffers a failure, or requires downtime, the *Standby Server* can quickly be brought on-line to host the database(s) from the QR Volumes you have created.

---

### STANDBY SERVER – DIFFERENT NAME FROM THE PRODUCTION SERVER

1. If necessary, reassociate all SQL scripts to the new machine name.
2. Start the SQL Services on the *Standby Server*.
3. Using SQL Enterprise Manager, attach the database(s) to the appropriate Instance. At this time, the database(s) will be online, and available to use.

---

### STANDBY SERVER – SAME NAME AS THE PRODUCTION SERVER

1. Remove the name of the source machine from Active Directory. This can be done by either renaming the source machine, or shutting it down and manually deleting the entry from *Active Directory Users and Computers*.
2. Rename the *Standby Server* to match the name of the *Production Server* that has been removed from Active Directory. If the machine name is still in Active Directory, allow time for Active Directory to Replicate throughout the domain.
3. Reboot the *Standby Server*; for the new machine name to take effect.
4. Start the SQL services on the *Standby Server*.
5. Using SQL Enterprise Manager, attach the database(s) to the appropriate instance. At this time, the database(s) will be online, and available to use.

[Back to Top](#)

---

# QR Disaster Recovery Solution for Building a Windows Standby Exchange Server in Quick Implementation Mode

---

## Overview

### Configuration

- Prepare the Production Server
- Prepare the Standby Server
- Configure the Quick Recovery Agent

### Bring the Standby Server online

- Standby Server Exchange Configuration

### Appendices

- Clean up after Quick Implementation for Exchange on a non-cluster machine
- 

## OVERVIEW

This document describes the procedure to create a *Standby Exchange Server* in the event that a *Production Exchange Server* is temporarily or permanently damaged.

These procedures are supported with Exchange 2000 and Exchange 2003 on Windows 2000, and Exchange 2003 on Windows 2003. It has been certified only with the Quick Recovery Agent and QSnap, using the QSnap snap engine for Windows 2000, or the VSS snap engine for Windows 2003. A familiarity with the functionality and configuration of the QR Agent is necessary in order to properly conduct this procedure.

This document contains four sections. First, the *Production Server* is configured. Second, the *Standby Server* is configured. Third, the Quick Recovery Volumes that will be used on the *Standby Server* are created. Fourth, the application is brought online.

### ADVANTAGES

- There is no downtime of the *Production Server* during Configuration.
- Recovery time does not include a system state restore.

### DISADVANTAGES

- Recovery time includes an Exchange installation and Service Pack installation.
- 

## CONFIGURATION

The following sections discuss preparing the *Production* and *Standby Servers* as well as configuring the QR Agent. The basic workflow is described below. For detailed instructions on installation and configuration options, see Quick Recovery Agent.

### PREPARE THE PRODUCTION SERVER

This procedure assumes that the *Production Server* has already been installed with Exchange 2000/2003 with the latest service packs or patches that may be needed, and that the users and mailboxes have already been created and configured.

1. Place the Exchange Transaction log location and the Exchange System path location on the same volume. The purpose of this step is to ensure that the `exchange.chk` file will be at the same point in time as the databases and logs when QR Volumes are created.

**NOTE:** Keep track of your installation selections and storage locations so that Exchange can be installed identically on the *Standby Server*.

2. Install the Quick Recovery Agent and QSnap on your production server. If you want to use VSS as your snap engine, install the VSS Enabler in addition to what's listed above. If necessary, install CommCell Console as a Stand-Alone Application. You don't need to install a MediaAgent or the Exchange Database `iDataAgent` on the *Production Server*.

### PREPARE THE STANDBY SERVER

Install the Quick Recovery Agent, QSnap, and MediaAgent software on the *Standby Server*. If necessary, install the CommCell Console as a Stand-Alone Application.

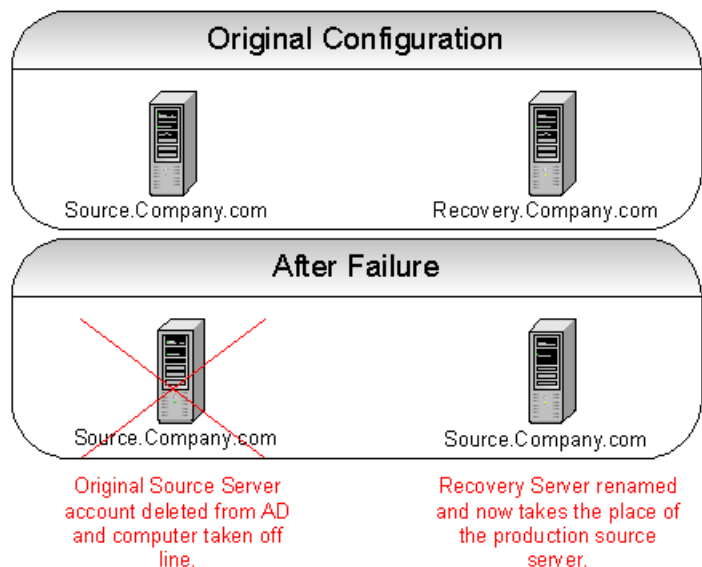
### CONFIGURE THE QUICK RECOVERY AGENT

1. Create a scratch volume pool that contains the destination volumes on the *Standby Server*.

2. Create a QR policy, set the appropriate snap engine type, use the *Standby Server's* LAN Copy manager, and associate your scratch volume pool.
3. Under your production server client, create a new QR subclient. Associate the QR policy created above with the new subclient.
4. When adding content to your subclient, use the **Add App** button to discover all the Exchange volumes on your *Production Server*.
5. When finished creating the subclient, right-click the subclient and select **Create QR Volumes**. Schedule an incremental update and use the **Advanced** button to map each source volume to its corresponding destination volume and mount point.
6. Create QR volume(s) from the *Production Server* to the *Standby Server*.

## BRING THE STANDBY SERVER ONLINE

The strategy for replacing the original *Production Server* with the *Standby Server* is illustrated below:



Due to dependencies that Exchange Server has within Active Directory, changes are required for Exchange to operate correctly on the *Standby Server*.

1. At this point assume that a problem has caused the *Production Server* to fail. Shut down the *Production Server* machine.
2. Reset the *Production Server* from Active Directory. To reset the *Production Server* from the *Active Directory Users and Computers* window, *Expand domain, Expand Computers*. Right-click on the *Production Server* and select **Reset Account**  
**NOTE:** Resetting a computer account breaks that computer's connection to the domain and requires it to rejoin the domain.
3. Take the *Standby Server* out of the domain and then rename the *Standby Server* to the Exchange Server (*Production Server*) name and reboot. Put the *Standby Server* back into the domain and reboot.
4. After the reboot, continue on to the *Standby Server* Exchange Configuration below.

## STANDBY SERVER EXCHANGE CONFIGURATION

1. Install Exchange on the *Standby Server* from the command line using the `/DisasterRecovery` command line switch. Select the same components, in Disaster Recovery mode, that were installed on the *Production Server*. (If this option does not exist, verify that the command line was entered correctly.)

Example: `Z:> setup.exe /DisasterRecovery`

**NOTE:** If Exchange won't allow you to run the setup with disaster recovery, you most likely have a permissions problem. First, try running the `setup.exe` command with `/DomainPrep`. After these changes have been made and replicated through the domain, you should be able to run the `setup /DisasterRecovery` command.

2. Do not reboot after Exchange has been installed.
3. Apply the same Service Packs and Patches that are on the original Server.

Example: `Z:> update.exe /DisasterRecovery`

**NOTE:** If Exchange won't allow you to run the update with disaster recovery, you most likely have a permissions problem. First, try running the `update.exe` command with `/DomainPrep`. After these changes have been made and replicated through the domain, you should be able to run the `update /DisasterRecovery` command.



4. After the installation, use Exchange System Manager to verify that each of the components selected are installed to the same location on the *Standby Server* as they are on the *Production Server*.

**NOTE:** You may see these informational messages, which do not apply to this configuration and may be disregarded:

- Please use Exchange Admin Snap-in to ensure that you have a valid Exchange Server Object for this server for which you are running setup in recovery mode.
- After setup has completed, please restore your databases from backup and then reboot your machine.

5. Reboot the *Standby Server*.

6. Verify that the Exchange Services have started without any errors and that all applicable Mailbox and Public Folder Stores have mounted.

**WARNING:** If a QR volume creation or incremental update is in the copy phase during a *Production Server* failure and has not completed, the data on the *Standby Server* will be incomplete and Exchange will not start.

**NOTE:** If any of the services fail to start the first time, you may need to manually start them.

7. If DNS is being used, edit the properties of the *Production Server* to point to the *Standby Server*.  
Or, change the IP address on the *Standby Server* to match the original *Production Server's* address.

Exchange is now operating on the *Standby Server*, which now takes the place of the *Production Server*.

## APPENDICES

### CLEAN UP AFTER QUICK IMPLEMENTATION FOR EXCHANGE ON A NON-CLUSTER MACHINE

1. Recovery Server
  - a. Delete the Outlook profile of the Exchange user.
  - b. Delete the user from **Active Directory User and Computers**.
  - c. Delete the Public and Mailbox Stores from the **Exchange System Manager**.
  - d. Delete the storage group from the **Exchange System Manager**.
  - e. To uninstall Microsoft Exchange, proceed as follows:
    - For Exchange 2003, see <http://support.microsoft.com/kb/833396/>.
    - For Exchange 2000, see <http://support.microsoft.com/?id=260378>.
  - f. When you have completed the procedures in the appropriate Microsoft article, if necessary, perform the following:
    - In the registry, delete the registry keys 65D9643D-06E8-47d6-865E-80F4CC9BB879 and 13F9F3AF-9463-4492-854A-191CCC441FDB from the following location:  
HKEY\_LOCAL\_MACHINE -> SOFTWARE -> Microsoft -> Windows -> CurrentVersion -> Uninstall
    - If you cannot delete three dlls (exchmem.dll, exsp.dll, pttrace.dll) in the Exchsrvr\bin folder, use procexp.exe to see which processes are using these dlls and kill them.
    - Delete the Exchsrvr folder.
  - g. Rename the Recovery Server back to the original name and reboot.
2. Source Server (only for Windows 2003)
  - a. The Source server has to be brought up in the Safe mode.
  - b. Set all the Exchange services to Manual and boot the system.
  - c. Even though it looks like your Source Server is in the domain, it is not. For this you have to take the Server out of the domain and reboot.
  - d. Log into the domain and reboot again.
  - e. Delete the Outlook profile of the Exchange user.
  - f. To uninstall Microsoft Exchange 2003, see <http://support.microsoft.com/kb/833396/>.
  - g. If you cannot delete three dlls (exchmem.dll, exsp.dll, pttrace.dll) in the Exchsrvr\bin folder, use procexp.exe to see which processes are using these dlls and kill them.
  - h. Delete the Exchsrvr folder.
3. Source Server (only for Windows 2000)

- a. Even though it looks like your Source Server is in the domain, it is not. For this you have to take the Server out of the domain and reboot.
- b. Log into the domain and reboot again.
- c. Delete the Outlook profile of the Exchange user.
- d. To uninstall Microsoft Exchange, proceed as follows:
  - o For Exchange 2003, see <http://support.microsoft.com/kb/833396/>.
  - o For Exchange 2000, see <http://support.microsoft.com/?id=260378>.
- e. In the registry, delete the registry keys 65D9643D-06E8-47d6-865E-80F4CC9BB879 and 13F9F3AF-9463-4492-854A-191CCC441FDB from the following location:  
HKEY\_LOCAL\_MACHINE -> SOFTWARE -> Microsoft -> Windows -> CurrentVersion -> Uninstall
- f. If you cannot delete three dlls (exchmem.dll, exsp.dll, pttrace.dll) in the Exchsrvr\bin folder, use procexp.exe to see which processes are using these dlls and kill them.
- g. Delete the Exchsrvr folder.

Back to Top

---

# QR Disaster Recovery Solution for Building a Windows Standby Exchange Server in Quick Recovery Mode

---

Overview

Configuration

- Prepare the Production Server
- Prepare the Standby Server
- Configure the Production and Standby Servers
- Configure the Quick Recovery Agent

Bring the Standby Server online

Appendices

- Clean up after Quick Recovery for Exchange on a non-cluster machine
- 

## OVERVIEW

This document describes the procedure necessary to create a *Standby Exchange Server* in the event that a *Production Exchange Server* is temporarily or permanently damaged.

These procedures are supported with Exchange 2000 and Exchange 2003 on Windows 2000, and Exchange 2003 on Windows 2003. It has been certified only with the Quick Recovery Agent, using the QSnap snap engine for Windows 2000, or the VSS snap engine for Windows 2003. A familiarity with the functionality and configuration of the QR Agent is necessary in order to properly conduct this procedure.

In the configuration stage, you will install Exchange in DisasterRecovery mode on the *Standby Server*. To use the *Standby Server*, you would reset the *Production Server* in Active Directory; then rename the *Standby Server* to the original production server name. Exchange functionality would then continue using the *Standby Server*.

This document contains four sections. First, the *Production Server* is configured. Second, the *Standby Server* is configured. Third, the Quick Recovery Volumes that will be used on the *Standby Server* are created. Fourth, the application is brought online.

### ADVANTAGES

- Very fast recovery time. Recovery time is limited to the length of time it takes to mount the volumes/stores and reconnect the mailboxes. Note that scripting these actions would speed the recovery.
- Recovery time does not include a system state restore, Exchange installation, service pack installation, etc.

### DISADVANTAGES

- Required scheduled down time of *Production Server* during configuration.
- 

## CONFIGURATION

The following sections discuss preparing the *Production* and *Standby Servers* as well as configuring the QR Agent. The basic workflow is described below. For detailed instructions on installation and configuration options, see Quick Recovery Agent.

### PREPARE THE PRODUCTION SERVER

This procedure assumes that the *Production Server* has already been installed with Exchange 2000/2003 with the latest service packs or patches that may be needed, and that the users and mailboxes have already been created and configured.

1. Place the Exchange Transaction log location and the Exchange System path location on the same volume. The purpose of this step is to ensure that the Exchange `chk` file will be at the same point in time as the databases and logs when QR Volumes are created.

**NOTE:** Keep track of your installation selections and storage locations so that Exchange can be installed identically on the *Standby Server*.


2. Install the Quick Recovery Agent and QSnap on your *Production Server*. If you want to use VSS as your snap engine, install the VSS Enabler in addition to what's listed above. If necessary, install CommCell Console as a Stand-Alone Application. You don't need to install a MediaAgent or the Exchange Database `iDataAgent` on the *Production Server*.

### PREPARE THE STANDBY SERVER

Install the Quick Recovery Agent, QSnap, and MediaAgent software on the *Standby Server*. If necessary, install the CommCell Console as a Stand-Alone Application.

## CONFIGURE THE PRODUCTION AND STANDBY SERVERS

1. On the *Production Server*:
  - a. Set the following services configuration, if they are not already set:
    - Microsoft Exchange Event should be stopped with the startup type set to **Manual**
    - Microsoft Exchange Site Replication Service should be stopped with the Startup Type set to **Disabled**
  - b. Stop the following Exchange Services, and set the Startup Type to **Manual**:
    - Microsoft Exchange IMAP4
    - Microsoft Exchange Information Store
    - Microsoft Exchange Management
    - Microsoft Exchange MTA Stacks
    - Microsoft Exchange POP3
    - Microsoft Exchange Routing Engine
    - Microsoft Exchange System Attendant
  - c. Rename the *Production Server* to a temporary name.  
Example: `<production server name>-temp`
  - d. Reboot the *Production Server*.
2. On the *Standby Server*:
  - a. Rename the *Standby Server* to the Exchange Server (*Production Server*).
  - b. Reboot the *Standby Server*.



- c. Create partitions and assign drive letters on the *Standby Server* to match those of the *Production Server* where the private and public store's databases and logs reside. (These partitions must be equal to, or greater than, the partition sizes of the *Production Server*.)
 

**NOTE:** The following components must be installed before installing Exchange: NNTP Services, and SMTP Services. Exchange 2003 also requires ASP .NET. These can be found on the Windows 2000/2003 install CD.
- d. Install Exchange on the *Standby Server* from the command line using the `/DisasterRecovery` switch. Select the same components, in DisasterRecovery mode, that were installed on the Production server. (If this option does not exist, verify that the command line was entered correctly).  
Example: `Z:> setup.exe /DisasterRecovery`

### NOTES:

- If Exchange won't allow you to run the update with disaster recovery, you most likely have a permissions problem. First, try running the `update.exe` command with `/DomainPrep`. After these changes have been made and replicated through the domain, you should be able to run the `update /DisasterRecovery` command.
- You may see these informational messages, which do not apply to this configuration and may be disregarded:
  - Please use Exchange Admin Snap-in to ensure that you have a valid Exchange Server Object for this server for which you are running setup in recovery mode.
  - After setup has completed, please restore your databases from backup and then reboot your machine.

- o When the installation reaches the post installation processing section, the program may hang while starting the System Attendant service. If this has happened, kill the `setup.exe` process, using the Windows Task Manager. Your ability to start services on the *Standby Server* will not be harmed. See Microsoft Knowledge Base article Q280432.
  - e. Apply the Service Packs and Patches that are on the Production server.  
Example: `Z:> update.exe /DisasterRecovery`  
**NOTE:** If Exchange won't allow you to run the update with disaster recovery, you most likely have a permissions problem. First, try running the `update.exe` command with `/DomainPrep`. After these changes have been made and replicated through the domain, you should be able to run the `update /DisasterRecovery` command.
  - f. On the *Standby Server*, stop the Exchange Services.
  - g. Set all of the Exchange Services to manual startup.
3. To enable the Exchange Services to start on the *Standby Server*, the Windows ADSI Edit tool must be installed. (See Microsoft KB article 325674). The tool can be installed on any server in the domain, but it is preferred to be installed on a Domain Controller. The ADSI tool is located on the Windows 2000/2003 Install CD. Once the tool is installed, you will need to set the appropriate permissions for the *Standby Server*. To start *ADSI Edit* click on **Start => Programs => Windows Support Tools => ADSI Edit**.
- a. In the ADSI Management console's left windowpane, expand the following:
    - o Configuration Container
    - o CN=Configuration
    - o CN=Services
    - o CN=Microsoft Exchange.
  - b. Right-click your organization [Domain Name] then click **Properties**. Click the **Security** tab:
    - o Verify that the *Standby Server* object is there. (*Production Server* name) If the server object is not listed, click **Add**, select the server object, click **Add**, and click **OK**.
    - o Highlight the server object and verify that the **Allow** check box for **Create all child objects** and **Delete all child objects** is selected. If they are not selected, click to select them, then click **OK**.
  - c. In the ADSI Management console's left windowpane, expand the following:
    - o CN=[Domain Name]
    - o CN=Administrative Groups
    - o CN=First Administrator Group
  - d. Right-click **CN=Servers** and select **Properties**. Click the **Security** tab:
    - o Verify that the *Standby Server* object is there. (*Production Server* name) If the server object is not listed, click **Add**, sort the objects by name, highlight the appropriate server object, click **Add** and click **OK**.
    - o Highlight the server object and click **Full Control** under Allow, then click **OK**.
  - e. Expand CN=Servers, right-click **CN=Servers** and select **Properties**. Click the **Security** tab:
    - o Verify that the *Standby Server* object is there. (*Production Server* name). If the server object is not listed, click **Add**, sort the objects by name, highlight the appropriate server object, click **Add** and click **OK**.
    - o Highlight the server object and click **Full Control** under Allow, then click **OK**.

Replicate these changes throughout the domain.

4. Manually Start the following Exchange Services on the *Standby Server*:
- o Microsoft Exchange IMAP4
  - o Microsoft Exchange Information Store
  - o Microsoft Exchange Management
  - o Microsoft Exchange MTA Stacks
  - o Microsoft Exchange POP3
  - o Microsoft Exchange Routing Engine
  - o Microsoft Exchange System Attendant
5. Open Exchange System Manager, then open **Organization Storage Group**, and right-click on the **Private Store Properties**. Click the **Database** tab and verify the directory paths for the `*.edb` and `*.stm` files are the same as the *Production Server*. If the paths do not match the *Production Server*, click **Browse**, and redirect the store to the correct path to match the *Production Server*. Do the same for the Public Store files.
6. Stop all of the Exchange Services on the *Standby Server*.
7. Rename the *Standby Server* back to its original name and reboot.

**NOTE:** Verify that the *Standby Server* has completely rebooted before continuing. If possible, check Active Directory to make sure the machine's name is

properly registered.

8. Rename the *Production Server* back to its original name and reboot.
9. Reset the Exchange services that were set previously to manual, back to automatic on the *Production Server*, and start these services.
10. Mount all mailbox and public folder stores that need to be mounted.
11. Verify that the stores will always mount at startup. (Right-click each store and select **Properties**. On the **Database** tab, disable **Do not mount this store at startup**.)

At this time, the *Production Server* (Exchange Server) is running again.

---

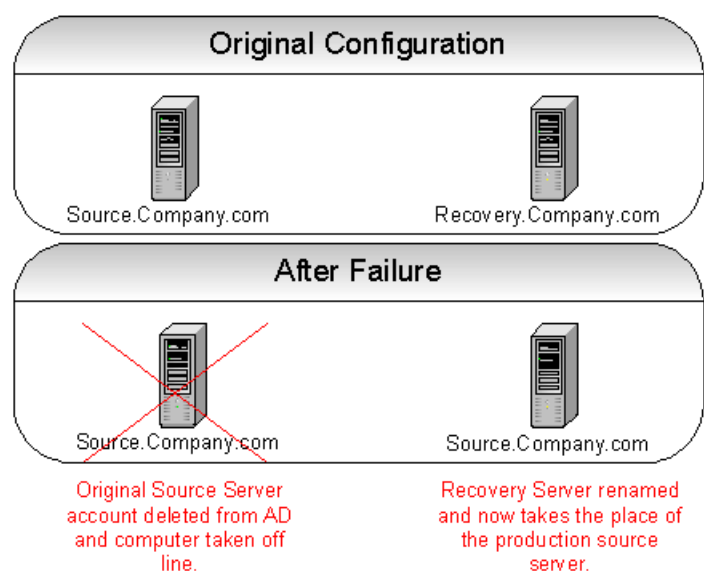
## CONFIGURE THE QUICK RECOVERY AGENT

1. Create a scratch volume pool that contains the destination volumes on the *Standby Server*.
2. Create a QR policy, set the appropriate snap engine type, use the *Standby Server's* LAN Copy manager, and associate your scratch volume pool.
3. Under your *Production Server* client, create a new QR subclient. Associate the QR policy created above with the new subclient.
4. When adding content to your subclient, use the **Add App** button to discover all the Exchange volumes on your *Production Server*.
5. When finished creating the subclient, right-click the subclient and select **Create QR Volumes**. Schedule an incremental update and use the **Advanced** button to map each source volume to its corresponding destination volume and mount point.
6. Create QR volume(s) from *Production Server* to the *Standby Server*.

---

## BRING THE STANDBY SERVER ONLINE

The strategy for replacing the original *Production Server* with the *Standby Server* is illustrated below:



Due to dependencies that Exchange Server has within Active Directory, changes are required for Exchange to operate correctly on the *Standby Server*.

1. At this point a problem has caused the *Production Server* to fail. Shut down the *Production Server* machine.
2. Reset the *Production Server* from Active Directory. To reset the *Production Server* from the Active Directory Users and Computers window, expand **Domain**, then expand **Computers**. Right click the *Production Server* and select **Reset Account**.  
**NOTE:** Resetting a computer account breaks that computer's connection to the domain and requires it to rejoin the domain.
3. Take the *Standby Server* out of the domain, rename the *Standby Server* to the Exchange Server (*Production Server*) name, and reboot. Put the *Standby Server* back into the domain and reboot.
4. After the reboot, set the following Exchange services to *Automatic* on the *Standby Server*, and start these services.
  - o Microsoft Exchange IMAP4
  - o Microsoft Exchange Information Store
  - o Microsoft Exchange Management

- o Microsoft Exchange MTA Stacks
- o Microsoft Exchange POP3
- o Microsoft Exchange Routing Engine
- o Microsoft Exchange System Attendant

**WARNING:** If a QR incremental update is in the copy phase during a *Production Server* failure and has not completed, the data on the *Standby Server* will be incomplete and Exchange will not start.

5. Mount all mailbox and public folder stores that need to be mounted.
6. If necessary, change the IP address on the *Standby Server* to match the original *Production Server's* address.

Exchange is now operating on the *Standby Server*, which now takes the place of the *Production Server*.

---

## APPENDICES

### CLEAN UP AFTER QUICK RECOVERY FOR EXCHANGE ON A NON-CLUSTER MACHINE

1. Recovery Server
  - a. Delete the Outlook profile of the Exchange user.
  - b. Delete the user from **Active Directory User and Computers**.
  - c. Delete the Public and Mailbox Stores from the **Exchange System Manager**.
  - d. Delete the storage group from the **Exchange System Manager**.
  - e. To uninstall Microsoft Exchange, proceed as follows:
    - o For Exchange 2003, see <http://support.microsoft.com/kb/833396/>.
    - o For Exchange 2000, see <http://support.microsoft.com/?id=260378>.
  - f. When you have completed the procedures in the appropriate Microsoft article, if necessary, perform the following:
    - o In the registry, delete the registry keys 65D9643D-06E8-47d6-865E-80F4CC9BB879 and 13F9F3AF-9463-4492-854A-191CCC441FDB from the following location:  
 HKEY\_LOCAL\_MACHINE -> SOFTWARE -> Microsoft -> Windows -> CurrentVersion -> Uninstall
    - o If you cannot delete three dlls (exchmem.dll, exsp.dll, pttrace.dll) in the Exchsrvr\bin folder, use procexp.exe to see which processes are using these dlls and kill them.
    - o Delete the Exchsrvr folder.
  - g. Rename the Recovery Server back to the original name and reboot.
2. Source Server (only for Windows 2003)
  - a. The Source server has to be brought up in the Safe mode.
  - b. Set all the Exchange services to Manual and boot the system.
  - c. Even though it looks like your Source Server is in the domain, it is not. For this you have to take the Server out of the domain and reboot.
  - d. Log into the domain and reboot again.
  - e. Delete the Outlook profile of the Exchange user.
  - f. To uninstall Microsoft Exchange 2003, see <http://support.microsoft.com/kb/833396/>.
  - g. If you cannot delete three dlls (exchmem.dll, exsp.dll, pttrace.dll) in the Exchsrvr\bin folder, use procexp.exe to see which processes are using these dlls and kill them.
  - h. Delete the Exchsrvr folder.
3. Source Server (only for Windows 2000)
  - a. Even though it looks like your Source Server is in the domain, it is not. For this you have to take the Server out of the domain and reboot.
  - b. Log into the domain and reboot again.
  - c. Delete the Outlook profile of the Exchange user.
  - d. To uninstall Microsoft Exchange, proceed as follows:
    - o For Exchange 2003, see <http://support.microsoft.com/kb/833396/>.
    - o For Exchange 2000, see <http://support.microsoft.com/?id=260378>.

- e. In the registry, delete the registry keys 65D9643D-06E8-47d6-865E-80F4CC9BB879 and 13F9F3AF-9463-4492-854A-191CCC441FDB from the following location:  
HKEY\_LOCAL\_MACHINE -> SOFTWARE -> Microsoft -> Windows -> CurrentVersion -> Uninstall
- f. If you cannot delete three dlls (exchmem.dll, exsp.dll, pttrace.dll) in the Exchsrvr\bin folder, use procexp.exe to see which processes are using these dlls and kill them.
- g. Delete the Exchsrvr folder.

[Back to Top](#)

---



# QR Disaster Recovery Solution for an Exchange Cluster Using a Standby Server

---

Overview

Configuration

- Prepare the Production Server
- Prepare the Standby Server
- Configure the Quick Recovery Agent

Bring the Standby Server online

---

## OVERVIEW

This document describes the procedure necessary to create a *Standby* Exchange cluster server in the event that a *Production* Exchange cluster server is temporarily or permanently damaged.

This procedure is supported with Exchange 2000 or Exchange 2003 on Windows 2000 on MSCS cluster or Exchange 2003 on Windows 2003 MSCS cluster. The cluster is assumed to be active/passive. This procedure has been certified with the Quick Recovery Agent and QSnap, using either the VSS snap engine for Windows 2003 or the QSnap snap engine for Windows 2000. A familiarity with the functionality and configuration of the Quick Recovery Agent is necessary to properly conduct this procedure.

The key to successfully recovering your Exchange data is proper configuration of the *Production* and *Standby Servers*. This ensures that the data can quickly and easily be brought online from the QR Volumes on the *Standby Server*.

The system administrator will prepare the *Standby* cluster server by installing Exchange on both physical nodes and using QR Agent to mirror the Exchange database and log volumes from the *Production Server*. In the event of a *Production Server* failure, the system administrator would create a new group on the *Standby* cluster using the *Production Server's* network name and IP. The QR Volume copies can then be added as physical disk resources and the Exchange System Attendant resource will be added to complete the transformation of your *Standby Server*.

### ADVANTAGES

- There is no downtime of the Production Server during configuration.
  - Recovery of the system state and Active Directory objects is not necessary.
  - Minimal time needed to bring up the Standby Server after a disaster on the Production Server.
- 

## CONFIGURATION

Refer to Quick Recovery Agent for details on installing and/or configuring QR Agent in a Clustered Environment

### PREPARE THE PRODUCTION SERVER

This section assumes that the Production cluster server has already been installed with Exchange 2000/2003 Server and with the latest service packs or patches that may be needed, and that users and mailboxes have already been created and configured.

1. Install QSnap and the Windows File System *iDataAgent* on the Production Cluster's *active physical node*, referred to as **Physical Node A**. If necessary, install the CommCell Console as a Stand-Alone Application.
2. Reboot **Physical Node A**. This will cause a failover, and **Physical Node B** will become the *active physical node*.
3. Install QSnap and the Windows File System *iDataAgent* on **Physical Node B**. If necessary, install the CommCell Console as a Stand-Alone Application.
4. Reboot **Physical Node B**. This will cause another failover, and **Physical Node A** will once again become the *active physical node*.
5. Install Quick Recovery Agent (and the VSS Enabler if you want to use VSS for your snaps) to the Cluster Server hosting the Exchange Server. QSnap is not needed on the Cluster Server. This installation should be done from the *active node* of the Production cluster, which at this time should again be **Physical Node A**. (reboot is not necessary after cluster server install).

#### NOTE:

For the <Software Installation folder>, select a volume that belongs to the cluster server hosting Exchange server but *not* the ones containing the data or logs.

6. Record your Exchange configuration on your *Production* cluster server. Specifically, record the Exchange transaction log location, Exchange system path location, database location, and streaming database location for each storage group that you would like to protect with this Disaster Recovery solution. Also, you will need to record your Exchange data directory, which is the location of your virtual Exchange server's installation. This information will be used

when setting up your destination volumes on your *Standby Server*.

7. From the Active Directory Users and Computers, verify that your Exchange server is a member of Exchange Domain Servers group. Right-click the server, select **Properties** and click the **Member of** tab. If the group is not listed add it.
8. Shut down the Virtual node resources of the *Production Server* if you are using the same Cluster Server name and IP address for the *Standby Server* virtual node that is used for the *Production Server* virtual name.

---

## PREPARE THE STANDBY SERVER

1. Based on Step 6 above, you will need to create the destination volumes, which will house your QR-copied Exchange data files. These destination volumes' drive letters should match those of their *Production Server* counterparts.

Here is an example assuming that you only want to protect the 'First Storage Group'. If on the *Production* cluster server –

- Exchange was installed on the Cluster Server to the **H:** drive
- The databases for First Storage Group were on the **G:** drive
- The transaction logs and system path were on the **F:** drive

You would need to create 3 volumes on the *Standby* cluster with the same drive letters.

QR Agent operates with Exchange on a storage group level, so you will only be copying the volumes which house your databases and transaction logs for a particular storage group. From the example above, that would mean you should copy your –

- **G:** drive on the *Production* cluster to the **G:** drive on your *Standby* cluster
- **F:** drive on the *Production* cluster to the **F:** drive on the *Standby* cluster.

You should make sure that your QR Volumes are the same size as, or larger than, their *Production Server* counterparts (this is just normal QR functionality). Although the contents of the data directory (the **H:** drive in the example) do not need to be copied to the *Standby* cluster in your QR operation, a volume with that drive letter needs to exist on your *Standby* cluster. When you install your cluster server on the *Standby* cluster during a disaster recovery, it will look for this data directory on your *Standby Server*.

### NOTE:

The destination volumes, which you create in this step, should **NOT** be assigned as resources to the cluster. QR Agent cannot copy to destination volumes which are cluster physical disk resources. Instead, you should just create and configure them on the *active physical node* of your *Standby* cluster. If the volumes are 'shared' (which is expected in a normal cluster setup) please remove the drive letters from these volumes on the *PASSIVE* nodes *OR* shut down all the passive nodes of the cluster. This ensures that they can only be written to on the Active physical node, thereby avoiding corruption issues. Once the QR operation is finished and a disaster occurs (assuring that QR operations will no longer occur), start the rest of the cluster nodes and reassign these volumes as physical disk resources to your *Standby* cluster.

2. Install Quick Recovery Agent, QSnap, and MediaAgent on the active physical node of your *Standby* cluster.

### NOTE:

You do not need to install any of these products on the passive node of your *Standby* cluster, or on the cluster server of your *Standby* cluster, in order for this procedure to work.

3. Install Microsoft Exchange on each physical nodes of the *Standby* cluster. Use the same path that is used on the *Production* cluster server (for example, if `C:\Program Files\Exchsrvr` is used on the *Production* cluster server for each Exchange installation, use that same path on your *Standby* cluster). Apply all necessary patches and service pack up to the level that Exchange on the *Production Server* is installed.

---

## CONFIGURE THE QUICK RECOVERY AGENT

1. Create a scratch volume pool: from the CommCell Console, create a new scratch volume pool which contains the volumes created/configured to house your database and transaction log data, as discussed in *Standby Server* configuration. Using the example above, this would be the **G:** drive and **F:** drive.
2. Create and configure a QR Policy:
  - a. Use VSS as your snap engine type for Windows 2003 clusters and QSnap as your snap engine type for Windows 2000 clusters.
  - b. Set the LAN Copy Manager of your *Standby Server's* active node as your copy manager.
  - c. Use the scratch volume pool just created in Step 1.
3. Create a new subclient on the cluster server of your *Production Server* and associate it with the QR Policy just created. Use the **Add App** button in the Subclient Properties Configuration tab to discover the Exchange volume for your storage group(s) that hosts your database files, transaction logs, and system path. From the example, this would be the **G:** drive and **F:** drive.

### NOTE:

In case **Add App** does not discover Exchange Server volumes, you might need to specify the name of the Exchange server (the name that appears in the Exchange system manager as your server). To do this, open the CommCell Console, expand your Production Server cluster server client, right-click **Quick Recovery Agent** and select **Properties**. In the Authentication tab, select **Exchange** and click **Edit**. Enter a valid user name and password for your Exchange server and then add your Exchange Server Name. This should fix any issues you may have in discovering Exchange on your cluster.

4. Set up a QR incremental update schedule so you can update your Exchange data on the Standby Server.

**NOTE:**

MAKE SURE THAT YOU CONFIGURE THE SCHEDULE SO THAT EACH VOLUME DRIVE LETTER ON THE PRODUCTION CLUSTER WILL BE COPIED TO THE SAME DRIVE LETTER ON THE STANDBY CLUSTER. From the example above, G: on the *Production Server* would copy to G: on the *Standby Server*, and F: on the *Production Server* would copy to F: on the *Standby Server*.

Refer to Quick Recovery Agent for details on configuring the Quick Recovery Agent.

---

## BRING THE STANDBY SERVER ONLINE

Due to dependencies that Exchange server has within Active Directory, changes are required to get Exchange operating correctly on the Standby cluster.

1. At this point of the procedure we assume that a problem has caused the *Production* cluster to fail.
2. Shut down all the physical nodes of the *Production* cluster server, one at a time.
3. Reset the Network Name for the *Production* cluster server from Active Directory Users and Computers. You can accomplish this by opening the **computers** directory and right-clicking your Production Exchange server and selecting **Reset Account**.
4. If you have shut down passive nodes on the *Standby Server*, start them one by one.
5. From Cluster Administrator of the *Standby* cluster, create a cluster server (a new group) using the *Production* Exchange server network name. Create the network name and IP address resources. You must use the same network name that was used in the *Production* cluster, but you can choose whether to use the same IP address or not. You may also need to enable Kerberos Authentication in the Parameters tab of your network name in order for the network name resource to start successfully.

**NOTE:**

If you use a different IP address, you will need to change some advanced settings in IIS in order to get things like Exchange's web server operational (HTTP and SMTP services), but your domain administrator should be able to handle this.

6. Add physical disk resources for both the QR Volumes and the Exchange data directory to the Cluster Server you just created. Using the example above, this would be your G:, F:, and H: drives. Do not add any dependencies to these physical disk resources.
7. Start all resources you have created in the new group. Verify they start successfully. Check MS Windows Event Viewer for error messages. Failover the group to the Passive node and back to the original Active node to verify services start.
8. From the Cluster Administrator, right-click the newly created group and select **New Resource** to create the Exchange System Attendant. This resource should have a dependency on your physical disk resources and the network name resource.
9. While creating the System Attendant Resource, verify the correct Exchange data directory is shown in the last step. In our example, this would be H:\EXCHSRVR.
10. After successful Exchange Resource creation, bring your Standby Server's cluster server online. Check for error messages and warnings in MS Windows Event Viewer.
11. At this time, you may open Outlook and verify your Exchange user's mailboxes. Your *Standby Server* is now your *Production Server*.

[Back to Top](#)

---

# QR Disaster Recovery Solution for Building a Standby SQL Server on a MS Cluster

---

Overview

Configuration

- Prepare the Standby Server
- Configure the Quick Recovery Agent

Bring the Standby Server online

---

## OVERVIEW

This document describes the procedure necessary to create a *Standby* SQL cluster server in the event that a *Production* SQL cluster server is temporarily or permanently damaged.

This procedure is supported with SQL Server, on either a Windows 2000 or Windows 2003 MSCS cluster. The cluster is assumed to be active/passive. This procedure has been certified with the Quick Recovery Agent and QSnap, using either the VSS snap engine for Windows 2003 or the QSnap snap engine for Windows 2000 / 2003. A familiarity with the functionality and configuration of the Quick Recovery Agent is necessary to properly conduct this procedure.

The key to successfully recovering your SQL data is proper configuration of the *Production* and *Standby Servers*. This ensures that the data can quickly and easily be brought online from the QR Volumes on the *Standby Server*.

## ADVANTAGES

- Minimal time needed to bring up the *Standby Server* after a disaster on the *Production Server*.
  - Very fast recovery time. Recovery time is limited to the time taken to attach the SQL databases.
- 

## CONFIGURATION

Refer to Quick Recovery Agent for details on installing and/or configuring QR Agent in a Clustered Environment

This section assumes that the *Production* cluster server has already been installed with SQL Server with the latest service packs or patches that may be needed, and that all user-defined SQL databases, tables, etc., have already been created.

1. Install QSnap and the Windows File System *iDataAgent* on the Production Cluster's *active physical node*, referred to as **Physical Node A**. If necessary, install the CommCell Console as a Stand-Alone Application.
2. Reboot **Physical Node A**. This will cause a failover, and **Physical Node B** will become the *active physical node*.
3. Install QSnap and the Windows File System *iDataAgent* on **Physical Node B**. If necessary, install the CommCell Console as a Stand-Alone Application.
4. Reboot **Physical Node B**. This will cause another failover, and **Physical Node A** will once again become the *active physical node*.
5. Install Quick Recovery Agent on the Virtual SQL Server. If you want to use VSS as your snap engine, then you should also install the VSS Enabler on the Virtual SQL Server. QSnap is not needed on the Cluster Server. This installation should be performed from the active node of the *Production* cluster, which at this time should again be **Physical Node A**. (reboot is **not** necessary after virtual node install).
6. Record your SQL configuration on your *Production* cluster server. Specifically, record the SQL transaction log location and database location (including all file groups) for each database that you would like to protect with this Disaster Recovery solution. Also, you will need to record the location of your virtual SQL server's installation. This information will be used when setting up your destination volumes on your *Standby Server*.
7. Shut down the Virtual node resources of the *Production Server* if you are using the same Cluster Server name and IP address for the *Standby Server* virtual node that is used for the *Production Server* virtual name.

---

## PREPARE THE STANDBY SERVER

1. Install the *Standby* cluster if it is not installed at this time. The cluster names and IP address must be different from the *Production* cluster.
2. Based on Step 6 above, you will need to create the destination volumes that will house your QR-copied SQL data files. These destination volumes' drive letters should match those of their *Production Server* counterparts. For example:

On the Production cluster server:

- SQL was installed on the virtual node to the H: drive
- the .mdf files were on the G: drive
- the .ldf files were on the F: drive

On the *Standby* cluster, create three volumes with the same drive letters. Ensure that your QR Volumes are the same size as, or larger than, their *Production* server counterparts (this is just normal QR functionality).

QR Agent operates with SQL on a database level, so you will only be copying the volumes that house your *.mdf*, *.ndf*, and *.ldf* files for a particular database. From the example above, that would mean you should copy:

- the **G:** drive on the *Production* cluster to the **G:** drive on your *Standby* cluster
- the **F:** drive on the *Production* cluster to the **F:** drive on the *Standby* cluster

Although the contents of the installation directory (the **H:** drive in the example) do not need to be copied to the *Standby* cluster in your QR operation, a volume with that drive letter needs to exist on your *Standby* cluster. When you install your virtual node on the *Standby* cluster during a disaster recovery, it will look for this data directory on your *Standby* server.

### NOTE:

The destination volumes that you will create above should **NOT** be assigned as resources to the cluster. QR Agent cannot copy to destination volumes that are cluster physical disk resources. Instead, you should just create and configure them on the active node of your *Standby* cluster. If the volumes are 'shared' (which is expected in a normal cluster setup) please remove the drive letters from these volumes on the PASSIVE node *OR* shut down all the passive nodes of the cluster. This ensures that they can only be written to on the *Production* node, thereby avoiding corruption issues. Once the QR operation is finished and a disaster occurs (assuring that QR operations will no longer occur), then you can reassign these volumes as physical disk resources to your *Standby* cluster.

3. Install Quick Recovery Agent, QSnap, Windows File System *iDataAgent*, and MediaAgent on the active node of your *Standby* cluster.

### NOTE:

You do **not** need to install any of these products on the passive node of your *Standby* cluster, or on the virtual node of your *Standby* cluster, in order for this procedure to work.

## STANDBY SERVER – DIFFERENT NAME FROM THE PRODUCTION SERVER

1. Install SQL and any Service Packs on the *Standby* cluster. Use the same installation path that is used on the *Production* cluster server. Also, use the same Virtual Name and IP address.
2. Use a different name and IP address from the one used for the *Production Server* Virtual node.

## STANDBY SERVER – SAME NAME AS THE PRODUCTION SERVER

1. Install SQL and any Service Packs on the *Standby* cluster. Use the same path that is used on the *Production* cluster server. Also, use the same Virtual Name and IP address.
2. Use the same Cluster Server name and IP address as the one used for the *Production Server* Virtual node.
3. Shut down all the resources of the SQL Group from the Cluster Administrator of the *Standby Server*.
4. Start all the resources of the SQL Group from the Cluster Administrator of the *Production Server*.

---

## CONFIGURE THE QUICK RECOVERY AGENT

1. Configure QR Agent in the CommCell Console so that your scratch volume pool contains the volumes created/configured to house your database and transaction log data, as discussed in Step 2 above. Using the example from Step 2, this would be the **G:** drive and **F:** drive.
2. Create and configure a QR Policy:
  - Use VSS as your snap engine type for Windows 2003 clusters and QSnap as your snap engine type for Windows 2000 clusters.
  - Set the LAN Copy Manager of your *Standby Server's* active node as your copy manager.
  - Use the scratch volume pool just created in Step 1.
3. Create a new subclient on the virtual node of your *Production Server* and associate it with the QR Policy just created. Use the **Add App** button in the Subclient Properties Configuration tab to discover the SQL volume for your storage group(s) that houses your database files and transaction logs. From the example, this would be the **G:** drive and **F:** drive.
4. Set up a QR incremental update schedule so you can update your SQL data on the *Standby Server*.

### NOTES:

- MAKE SURE THAT YOU CONFIGURE THE SCHEDULE SO THAT EACH VOLUME DRIVE LETTER ON THE *PRODUCTION* CLUSTER WILL BE COPIED TO THE SAME DRIVE LETTER ON THE *STANDBY* CLUSTER. From the example above, **G:** on the *Production Server* would copy to **G:** on the *Standby Server*, and **F:** on the *Production Server* would copy to **F:** on the *Standby Server*.
- In addition, verify the binding order of your network Resources. The Public Address must be first.

Refer to Quick Recovery Agent for details on configuring the Quick Recovery Agent.

---

## BRING THE STANDBY SERVER ONLINE

1. At this point of the procedure we assume that a problem has caused the *Production* cluster to fail.
2. Shut down all the physical nodes of the *Production* cluster server, one at a time.
3. If you have shut down the passive nodes on the *Standby Server*, start them one by one.
4. For the *Standby* cluster's SQL cluster server, add physical disk resources for both the QR Volumes (the SQL data directory should already be a physical disk resource on the *Standby* SQL cluster server).  
Using the example above, this would be your **G:** and **F:** drives. Don't add any dependencies to these physical disk resources.
5. Add these physical disk resources as dependencies to the 'SQL Server' resource on your *Standby* SQL cluster server.
6. After successful SQL Resource creation, bring your *Standby Server's* virtual node online.
7. At this time, you may open Enterprise Manager and attach your *Production* SQL server's database using the QR Volumes on your *Standby* SQL server.  
Select the `.mdf` file from one of your QR Volumes in order to attach. Verify your SQL database, file groups, and tables have been preserved.

[Back to Top](#)

---