Features - ContinuousDataReplicator

TABLE OF CONTENTS

Overview - ContinuousDataReplicator

Choose from the following topics:

Introduction

- Data Replication
- Snapshots
- Application Awareness
- Recovery Points

Supported Data Types

- Windows File Systems
- Unix File Systems

Tree Levels in ContinuousDataReplicator

License Requirement

INTRODUCTION

ContinuousDataReplicator (CDR) replicates data from a source computer to a destination computer. This is done in near real-time by logging all file write activity to a replication log in the source computer, including new files and changes to existing files. These replication logs are transferred to the destination computer and replayed, ensuring that the destination remains a *nearly* real-time replica of the source.

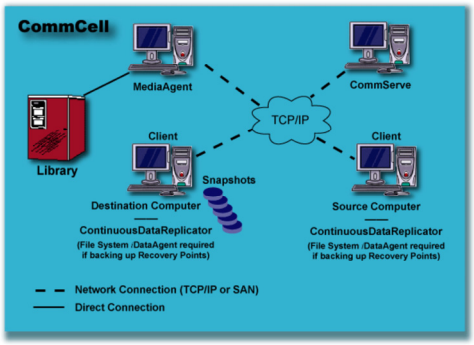

CDR provides protection for all types of data including file system and application data.

CDR provides easy-to-use wizards and policies to simplify the basic configuration, operations, and monitoring of data replication activities. In addition, a separate Data Replication Monitor is provided in the CommCell Console to track all replication activities. The Data Replication Monitor also provides specific reporting and activity history for additional monitoring.

CDR is just one of the several Agents that provide data protection and recovery support for different kinds of data in what is often a complex environment. The sample figure shows an example of a simple CommCell configuration using CDR.

CDR can also be configured and used for backing up computers in remote locations. See Disk Library Replication for more information.

CDR provides the following important features:

DATA REPLICATION

Data Replication is the process of copying the specified, file-level content from one computer, the source computer, to another, the destination computer. There are several scenarios for data replication can be configured. They are:

- Direct replication - replication from a source to a destination computer.
- Fan-In configuration - replication from multiple source computers to a single destination computer.
- Fan-Out configuration - replication from one source computer to multiple destination computers.

For more information on these scenarios, see Supported Configuration.

SNAPSHOTS

Snapshots represent a point-in-time of the data that can be used for various data protection operations. A snapshot is essentially an instantaneous set of pointers to the original data (sometimes referred to as a logical view) as it was at a given point-in-time. When the original data is changed, the pointers will trigger a copy of the original data block; this maintains the snapshot, allowing data protection operations to proceed without interruption.

CDR utilizes snapshots for certain phases of its operation, both on a source computer as well as a destination. For more information, see Snapshot Engines - Support. The snapshots supported for different operating systems are as follows:

- CDR on Windows supports:
  - VSS
  - ONTAP
- CDR on Unix supports:
  - File System Snapshot

- QSnap (Linux and Solaris)
- ONTAP

For general information about snapshots, see Snapshots.

APPLICATION AWARENESS

CDR supports the replication of data associated with applications, such as SQL, Exchange, etc. For a complete listing of applications supported for each operating system, see ContinuousDataReplicator - Application Support.

For supported applications, CDR will automatically discover directories that should be replicated (logs, databases, etc.) to protect the application's data. For more information, see Application Integration.

RECOVERY POINTS

CDR provides the ability to create Recovery Points which consists of snapshots of data that preserve a point-in-time for recovery operations. This provides high availability of protected data. For more information, see Recovery Points.

Consistent Recovery Points which are point-in-time snapshot of application data in a consistent state in the source computer (the application server) can also be created using CDR. This ensures that the consistent state of the application is preserved which in turn can be restored to the specific point-in-time. For more information see Consistent Recovery Points.

Both Recovery Points and Consistent Recovery Points can be mounted, shared as a volume on the network, or recovered using Copyback. After a Recovery Point has been created, the data can also be backed up, using the appropriate File System *i*DataAgent. This provides additional protection for your data, as well as additional options for recovery of your data. For more information, see Recovery Scenarios.

SUPPORTED DATA TYPES

CDR protects the following data types. Note that Windows data can only be replicated to another Windows computer and UNIX data can only be replicated to another UNIX computer.

WINDOWS FILE SYSTEMS

- File Allocation Table (FAT) file systems (Disk Library Replication only)
- New Technology File Systems (NTFS)
- Access Control Lists (ACL)
- Unicode files
- Sparse files (except where the destination is a NetApp filer)
- Single Instance Storage (for Windows Storage Server only)
- Mount Points
- Encrypted files (except where the destination is a NetApp filer) - the key is not replicated, only the file itself (Also see Replicating encrypted data.)
- Shared Volumes
- Compressed Data (except where the destination is a NetApp filer) - including a compressed file contained in a compressed folder, any folder (not compressed) that contains two or more compressed files, or a compressed drive.
- Dynamic volumes are supported as Replication Pair content
- Software and Hardware RAID storage
- Migration files (stub only)
- iSCSI LUNs on destination computer
    - iSCSI dynamic disks are not supported.
    - iSCSI LUNs with multiple partitions are not supported.
- When replicating a root volume (e.g., `C:\`), paging files, the System Volume folder, and the hibernation files will be filtered out automatically from Replication Pair content, and will not be replicated to the destination machine.
- It is recommended that you use recovery points on the destination for examining consistency of data on a destination computer, as comparing live data on source and destination may result in the files not matching.

UNIX FILE SYSTEMS

- Logical volumes for AIX, Linux and Solaris; Linux and Solaris supports file systems on physical disks configured as CXBF devices
- Access Control Lists (ACL)
- Unicode files; files with non-ASCII characters in their name

- Root File Systems on source
- Sparse files

  During the Baselining phase, CDR on Unix will transfer sparse files as regular files. However, during the Replication phase, sparse files will be replicated as sparse files, and regions of the file that do not require disk space will be unallocated.

- Migration files (stub only)
- Linux:
  - Extent 2 File System (ext2)
  - Extent 3 File System (ext3)
  - Extent 4 File System (ext4)
  - Reiser File System (reiserfs)
  - VERITAS File System (VxFS) on volumes created by VERITAS Volume Manager
  - XFS File System (xfs)
  - Global File System 2 (GFS2)
- AIX:
  - Enhanced Journaled File System (JFS2)
  - VERITAS File System (VxFS)
  - Journaled File System (JFS) on source.

    The memory mapped files are not supported for AIX and Solaris.

- Solaris:
  - Zettabyte File System (ZFS)
  - Unix File System (UFS)
  - VERITAS File System (VxFS)

    If you are using UFS and VxFS File System, you need to install QSnap on the destination computer to create snapshots for Recovery Points.

- HP-UX:
  - VERITAS File System (VxFS)

    - Recovery Points and Consistent Recovery Points on destination are supported with VERITAS Volume Manager.
    - Snapshots on source are supported with VERITAS Volume Manager and OnlineJFS.

TREE LEVELS IN CONTINUOUSDATAREPLICATOR

When CDR is installed, Replication Sets and Replication Pairs are not created by default, you must create Replication Sets and Replication Pairs in the CommCell Browser.

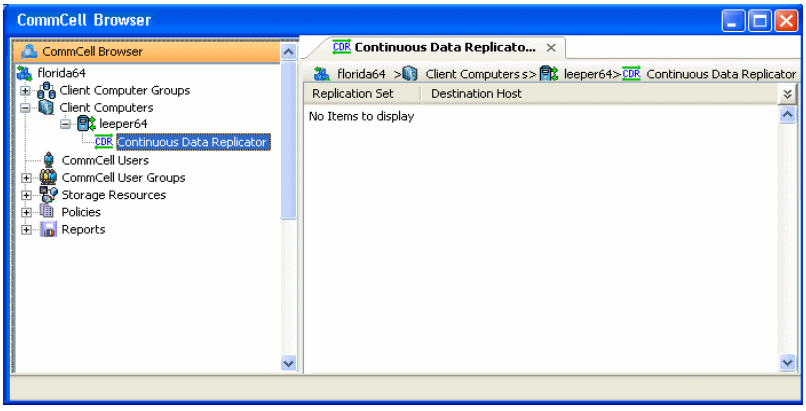

`leeper64:` Client `H:\data: I:\data33232`

`Continuous Data Replicator`: Agent `(user defined)`: Replication Pair

`Set1 (user defined)`: Replication Set

LICENSE REQUIREMENT

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.

The ContinuousDataReplicator requires the following licenses:

- **ContinuousDataReplicator** - This license is consumed during the installation of the software and each installed instance uses a license.
- **Recovery Points** license is required to use Recovery Point feature.

Back to Top

System Requirements - ContinuousDataReplicator

System Requirements | Snapshot Support | Supported Features

The following requirements are for the ContinuousDataReplicator:

OPERATING SYSTEM		PROCESSOR
AIX	AIX 7.1	Power PC (Includes IBM System p)
	AIX 6.1 64-bit	Power PC (Includes IBM System p)
	AIX 5.3 32-bit and 64-bit with technology level 7 (or higher) and runtime library xlC.rte 8.0.0.0 or higher Note that all AIX platforms support Logical Partitions (LPAR).	Power PC (Includes IBM System p)
HP-UX	HP-UX 11i v3 (11.31)	Itanium
	HP-UX 11i v3 (11.31)	PA-RISC
	HP-UX 11i v2 (11.23)	PA-RISC
	HP-UX 11i v2 (11.23)	Itanium
LINUX	**DEBIAN**	
	Debian 5.x with kernel 2.6.26-2	Intel Pentium or compatible minimum required
	Debian 5.x with kernel 2.6.26.19-2	x64
	RED HAT ENTERPRISE LINUX	
	Red Hat Enterprise Linux 6 Advanced Platform with kernel 2.6.32-71	Intel Pentium or compatible minimum required
	Red Hat Enterprise Linux 6 Advanced Platform with kernel 2.6.32-71	x64
	Red Hat Enterprise Linux 6 Advanced Platform with kernel 2.6.32-131 (Update 1)	x64
	Red Hat Enterprise Linux 5 Advanced Platform with kernel 2.6.18-92 (Update 2)	Intel Pentium or compatible minimum required
	Red Hat Enterprise Linux 5 Advanced Platform with kernel 2.6.18-92 (Update 2)	x64
	Red Hat Enterprise Linux 5 Advanced Platform with kernel 2.6.18-8	Intel Pentium or compatible minimum required
	Red Hat Enterprise Linux 5 Advanced Platform with kernel 2.6.18-8	x64
	Red Hat Enterprise Linux 5 Advanced Platform with kernel 2.6.18-53 (Update 1)	x64
	Red Hat Enterprise Linux 5 Advanced Platform with kernel 2.6.18-53 (Update 1)	Intel Pentium or compatible minimum required
	Red Hat Enterprise Linux 5 Advanced Platform with kernel 2.6.18-238 (Update 6)	Intel Pentium or compatible minimum required
	Red Hat Enterprise Linux 5 Advanced Platform with kernel 2.6.18-238 (Update 6)	x64
	Red Hat Enterprise Linux 5 Advanced Platform with kernel 2.6.18-194 (Update 5)	x64
	Red Hat Enterprise Linux 5 Advanced Platform with kernel 2.6.18-194 (Update 5)	Intel Pentium or compatible minimum required
	Red Hat Enterprise Linux 5 Advanced Platform with kernel 2.6.18-164 (Update 4)	Intel Pentium or compatible minimum required
	Red Hat Enterprise Linux 5 Advanced Platform with kernel 2.6.18-164 (Update 4)	x64
	Red Hat Enterprise Linux 5 Advanced Platform with kernel 2.6.18-128 (Update 3)	x64
	Red Hat Enterprise Linux 5 Advanced Platform with kernel 2.6.18-128 (Update	Itanium

	3)	
	Red Hat Enterprise Linux 5 Advanced Platform with kernel 2.6.18-128 (Update 3)	Intel Pentium or compatible minimum required
	SUSE LINUX (SLES)	
	SuSE Linux 11 Enterprise Server with kernel 2.6.32.12-0.7 (Update 1)	Intel Pentium or compatible minimum required
	SuSE Linux 11 Enterprise Server with kernel 2.6.32.12-0.7 (Update 1)	x64
	SuSE Linux 11 Enterprise Server with kernel 2.6.27.19-5	Intel Pentium or compatible minimum required
	SuSE Linux 11 Enterprise Server with kernel 2.6.27.19-5	x64
	SuSE Linux 10 Enterprise Server with kernel 2.6.16.60-0.54.5 (Update 3)	Intel Pentium or compatible minimum required
	SuSE Linux 10 Enterprise Server with kernel 2.6.16.60-0.34 (Update 2)	Intel Pentium or compatible minimum required
	SuSE Linux 10 Enterprise Server with kernel 2.6.16.60-0.21 (Update 2)	x64
	SuSE Linux 10 Enterprise Server with kernel 2.6.16.60-0.21 (Update 2)	Intel Pentium or compatible minimum required
	SuSE Linux 10 Enterprise Server with kernel 2.6.16.46-0.12 (Update 1)	Intel Pentium or compatible minimum required
	SuSE Linux 10 Enterprise Server with kernel 2.6.16.46-0.12 (Update 1)	x64
SOLARIS	Solaris 9	Sparc5 or higher recommended
	Solaris 10.x with a minimum of SunOS (Sparc) Patch 119963-14	Sparc5 or higher recommended
	Solaris 10.x (6/06 and higher)	Sparc5 or higher recommended
	Solaris 10.x (6/06 and higher)	x64
WINDOWS	**WINDOWS 2008**	
	Microsoft Windows Server 2008 32-bit and x64 Editions	All Windows-compatible processors supported
	WINDOWS VISTA	
	Microsoft Windows Vista 32-bit and x64 Editions	All Windows-compatible processors supported
WINDOWS SERVER 2003	Microsoft Windows Server 2003 32-bit and x64 Editions with a minimum of Service Pack 2	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

CDR ON WINDOWS:

260 MB minimum of hard disk space for software

50 MB of additional hard disk space for log file growth

720 MB of temp space required for install or upgrade (where the temp folder resides)

- 1 GB minimum of additional hard disk space on the source computer for replication log file growth; more is recommended
- 1 GB of additional hard disk space on the destination computer for replication log file growth for each source computer (e.g., if five source computers are configured to use the same destination computer, then 5 GB of additional hard drive space is required on the destination computer)

CDR ON UNIX:

260 MB minimum of hard disk space for software

- 3 GB minimum of additional hard disk space on the source computer for replication log file growth; more is recommended. Maximum allowed is 80% of total space on the file system.

  Refer to Data Replication for important considerations when allocating replication log file space.

MEMORY

AIX AND LINUX

16 MB RAM minimum required beyond the requirements of the operating system and running applications

Swap space = 2*RAM size

WINDOWS

512 MB RAM minimum required, of which a minimum of 64 MB RAM is required beyond the requirements of the operating system and running applications for successful ContinuousDataReplicator operations; 1 GB RAM recommended.

SOLARIS ZONES/CONTAINERS SUPPORT

The ContinuousDataReplicator should be installed on the global zone. Installation on non-global zones is not supported.

On the source computer, the ContinuousDataReplicator can be used to replicate data from any non-global zone.

On the destination computer, data can be replicated to a non-global zone through a exporting device from the global zone. See Data Replication on Non-Global Zones for more information.

For a comprehensive list of supported components, see Unix Virtualization.

AIX LPAR/WPAR SUPPORT

Data protection on Logical Partitioning (LPAR) and Workload Partitioning (WPAR) is supported.

PERIPHERALS

DVD-ROM drive

Network Interface Card

MISCELLANEOUS

NETWORK

TCP/IP Services configured on the computer.

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

`<directory>/<file_name>.te`
 	
where:
 	
`<directory>` is `/usr/share/selinux/devel`
 	
`<file_name>` 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: `backup_IDA.te`
	The content of the file should be as follows:
	
	policy_module(<name>,<version>)
	
	`##############################`
	
	where:
	
	`<name>` is the name of the policy module. You can give any unique name to the policy module, such as a process or application name.
	
	`<version>` 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.
	
	`policy_module(backup_IDA,1.0.0)`

`<directory>/<file_name>.fc`
 	
where:
 	
`<directory>` is `/usr/share/selinux/devel`
 	
`<file_name>` 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
	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 `/var/log/messages`. Also, if required, add it to the following list of files.
	
	`/opt/<software installation directory>/Base/libCTreeWrapper.so -- gen_context (system_u:object_r:texrel_shlib_t,s0)`
	
	`/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`

backup_IDA application, you can use the following file name:
`backup_IDA.fc`
	`(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.

NOTES ON WINDOWS INSTALLATION

Microsoft Windows XP is not supported as both source and destination for ContinuousDataReplicator.

When installing ContinuousDataReplicator client on a Microsoft Windows x64 platform, you must install Microsoft .NET Framework 2.0.

NOTES ON LINUX INSTALLATION

A compiled version of the CDR driver is supplied only for the listed versions and kernels. If you perform online updates of Linux, you may encounter a situation where the supplied CDR driver will not load after a reboot because the kernel has been updated to a new version.

If you are installing the software on a computer running Red Hat Linux 5 or above, the file system on the computer must be compiled as a kernel module.

For example: computers running Red Hat Linux 5 with the `ext2` file system will not support the ContinuousDataReplicator software, as the ext2 file system is built into the kernel. Conversely, computers running Red Hat Linux 5 with the `ext3` file system compiled as a kernel module will support the ContinuousDataReplicator software.

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 ContinuousDataReplicator

TABLE OF CONTENTS

INSTALL REQUIREMENTS

The following procedure describes the steps involved in installing ContinuousDataReplicator.

ContinuousDataReplicator is installed on at least two machines; the server from which you will replicate data (source computer) and the computer to which you will replicate the data (destination computer.) You may choose to perform additional installations, based on the Supported Configuration used in your particular environment.

Verify that the computers in which you are installing the software satisfy the minimum requirements specified in System Requirements - ContinuousDataReplicator.

Review the following Install Requirements before installing the software:

GENERAL

- Review Install Considerations before installing the software.
- If any of the computers in which you are installing this software have multiple Network Interface Cards (NIC) you must configure them so that the source and destination computers can communicate for replication activities. For more information, see Data Interface Pairs.
- Agents should only be installed after the CommServe and at least one MediaAgent have already 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.
- This version of the software is intended to be installed in a CommCell where the CommServe and MediaAgent(s) version is 9.0.0.
- 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 for the Agent.
- Verify that the software installation disc is appropriate to the operating system of the computer on which the software is being installed.
  Make sure that you have the latest software installation disc before you start to install the software. If you are not sure, contact your software provider.

BEFORE YOU BEGIN

- Log on to the client as the local Administrator or as a member of the Administrators group on that computer.

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.

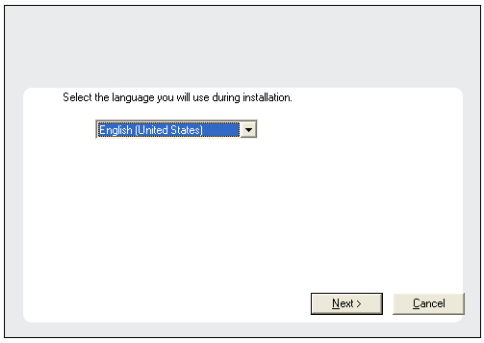

3. Select the option to install software on this computer.

NOTES

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

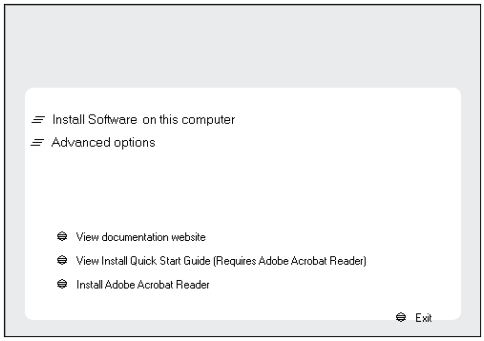

4. Read the license agreement, then select **I accept the terms in the license agreement**.

Click **Next** to continue.

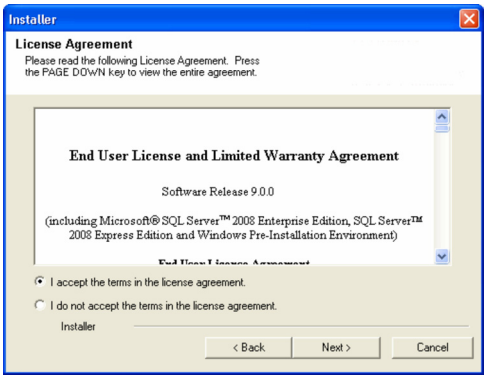

SELECT COMPONENTS FOR INSTALLATION

5. 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 ContinuousDataReplicator, expand the `Client Modules` folder and the `ContinuousDataReplicator` folder and select the following:

- `ContinuousDataReplicator`

To install VSS Provider for CDR, select the following:

- `VSS Provider`

For more information on VSS Provider for CDR, see VSS for ContinuousDataReplicator.

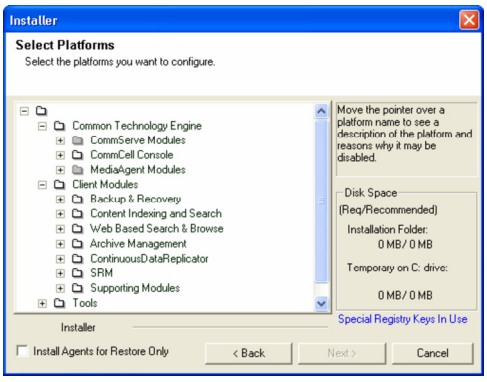

CONFIGURATION OF OTHER INSTALLATION OPTIONS

6. 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.

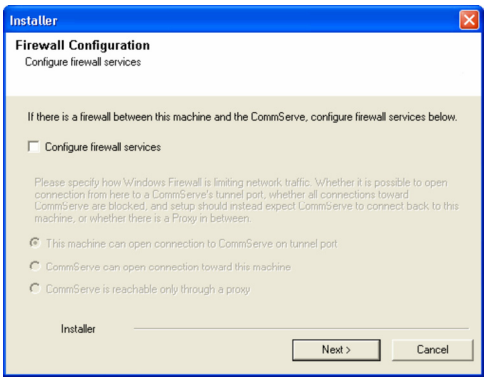

7. Enter the fully qualified domain name of the CommServe Host Name. This should be TCP/IP network name. e.g., computer.company.com.

   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.

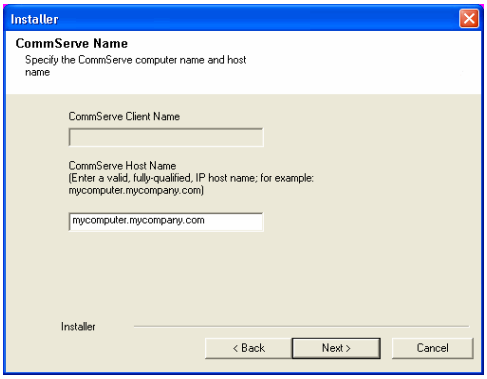

8. Enter the **username** and **password** associated with an external domain user account or a CommCell user account to authorize the installation of this agent.

   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.

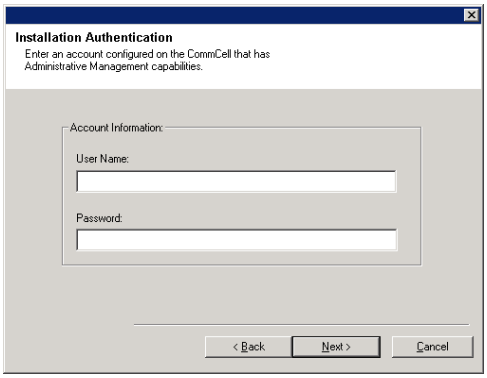

9. Enter the following:

   - 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.

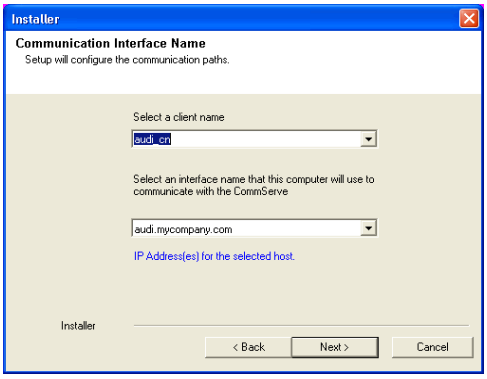

10. Select **Add programs to the Windows Firewall Exclusion List**, if you wish to add CommCell programs and services to the Windows Firewall Exclusion List.

    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.

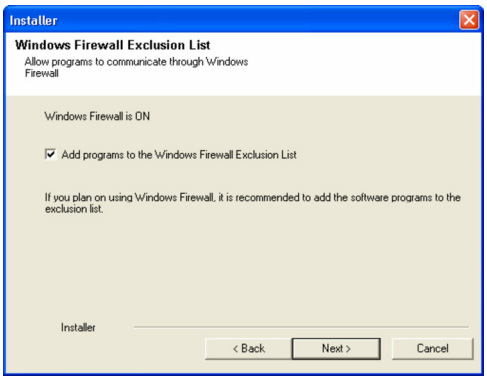

11. Specify the location where you want to install the software.

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.

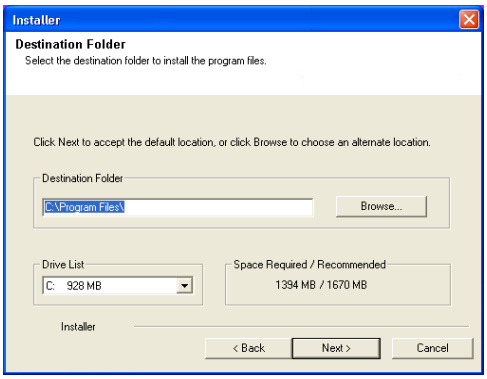

CLIENT GROUP SELECTION

12. 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.

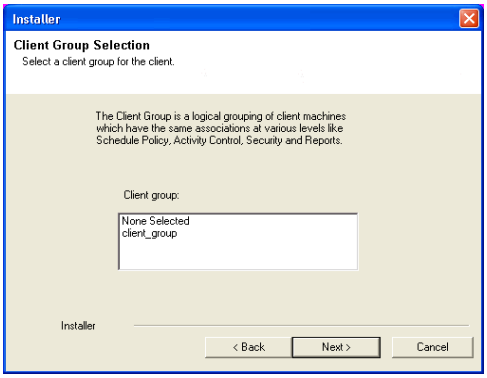

SCHEDULE AUTOMATIC UPDATE

13. 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.

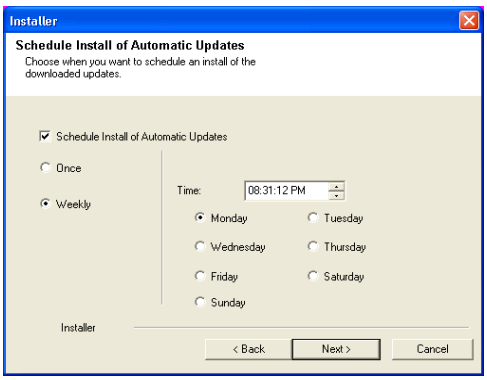

REPLICATION LOGS LOCATION

14. Specify a path for the Replication Logs directory.

Click **Browse** to change directories.

Click **Next** to continue.

NOTES

- There are several considerations for log file space and location; refer to Replication Logs when deciding on a suitable location for Replication Logs.

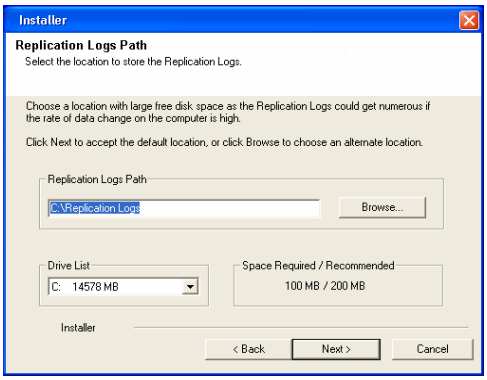

VERIFY SUMMARY OF INSTALL OPTIONS

15. 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.

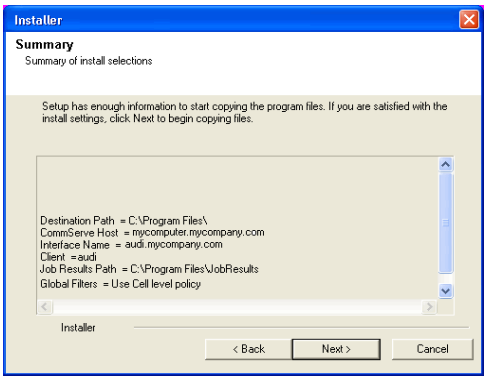

16. The System Reboot message may be displayed. If so, select one of the following:
    - **Reboot Now**
      If this option is displayed without the **Skip Reboot** option, the install program has found files required by the software that are in use and need to be replaced. If **Reboot Now** is displayed 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**.

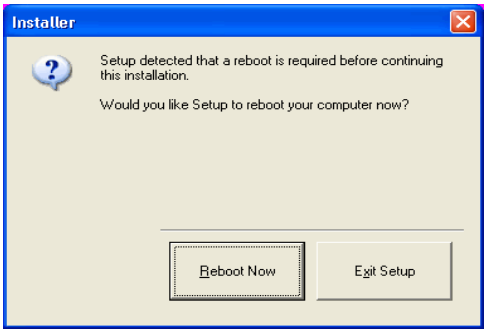

17. Setup reminds you that the computer must be restarted, after the installation completes, before you can use this Agent.

Click **OK** to continue.

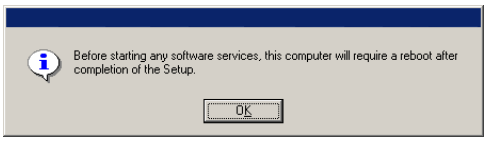

18. Click **Next** to continue.

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.

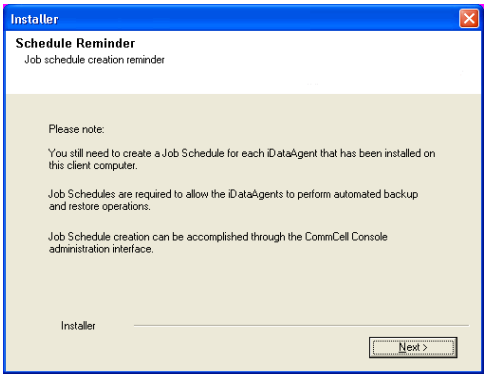

SETUP COMPLETE

19. 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.

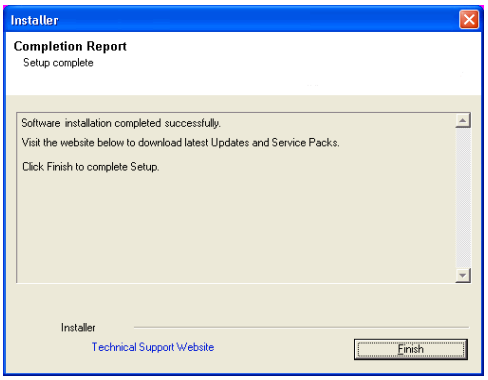

POST-INSTALL CONSIDERATIONS

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 Driver for the ContinuousDataReplicator - Clustered Environment - Physical Node

TABLE OF CONTENTS

INSTALL REQUIREMENTS

The following procedure describes the steps involved in installing the Driver for ContinuousDataReplicator. Before you can install ContinuousDataReplicator in the cluster group in a clustered environment, you must first install the Driver for ContinuousDataReplicator on all physical nodes of the cluster.

ContinuousDataReplicator is installed on at least two machines; the server *from* which you will replicate data (source computer) and the computer to which you will replicate the data (destination computer.) You may choose to perform additional installations, based on the Supported Configuration used in your particular environment.

Verify that the computers in which you are installing the software satisfy the minimum requirements specified in System Requirements - ContinuousDataReplicator.

Review the following Install Requirements before installing the software:

GENERAL

- If any of the computers in which you are installing this software have multiple Network Interface Cards (NIC) you must configure them so that the source and destination computers can communicate for replication activities. For more information, see Data Interface Pairs.
- Agents should only be installed after the CommServe and at least one MediaAgent have already 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.
- This version of the software is intended to be installed in a CommCell where the CommServe and MediaAgent(s) version is 9.0.0.
- 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 for the Agent.
- Verify that the software installation disc is appropriate to the operating system of the computer on which the software is being installed.
  Make sure that you have the latest software installation disc before you start to install the software. If you are not sure, contact your software provider.

CLUSTER

- 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.)

AGENT SPECIFIC

- In addition to a license for ContinuousDataReplicator, you must have an available license on the CommServe for any snapshot engines and copy managers you plan to use.

BEFORE YOU BEGIN

- 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.

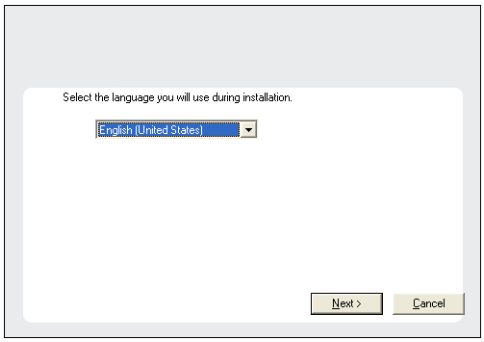

3. Select the option to install software on this computer.

   NOTES

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

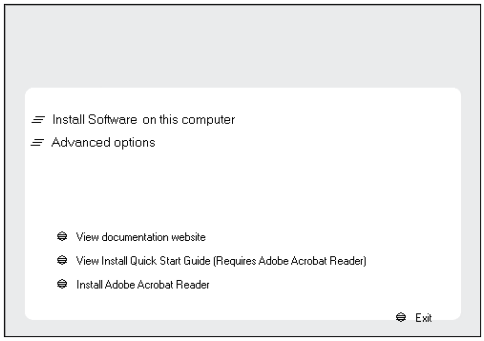

4. Read the license agreement, then select **I accept the terms in the license agreement**.

   Click **Next** to continue.

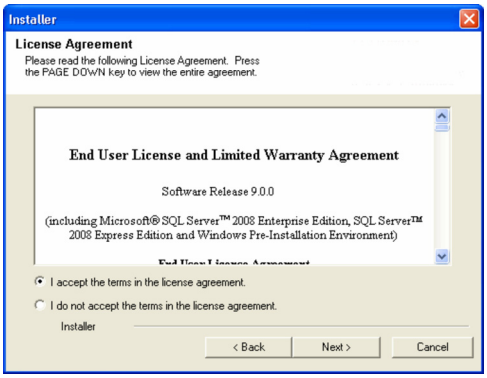

CLUSTER SELECTION

5. Select **Install on the Physical Node**.

   Click **Next** to continue.

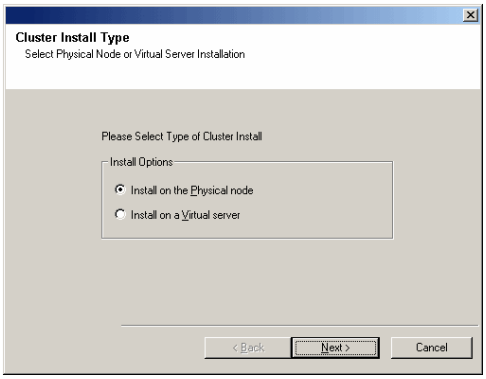

SELECT COMPONENTS FOR INSTALLATION

6. Select the component(s) to install.

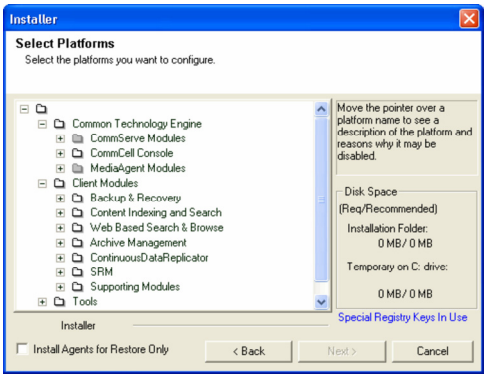

NOTES

- 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 ContinuousDataReplicator, expand the `Client Modules` folder and the `ContinuousDataReplicator` folder and select the following:

- `Driver for ContinuousDataReplicator`

To install VSS Provider for CDR, select the following:

- `VSS Provider`

VSS Provider for CDR must be installed on all physical nodes of a cluster. You cannot install it on the virtual node. For more information on VSS Provider for CDR, see VSS for ContinuousDataReplicator.

CONFIGURATION OF OTHER INSTALLATION OPTIONS

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

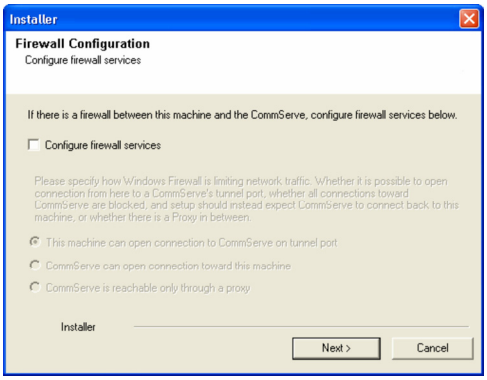

   For firewall options and configuration instructions, see Firewall Configuration and continue with the installation.

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

8. 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.

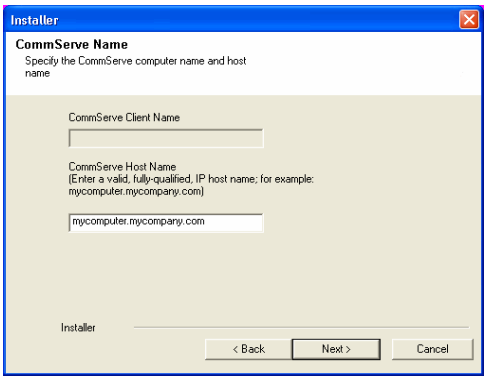

9. Enter the **username** and **password** associated with an external domain user account or a CommCell user account to authorize the installation of this agent.

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.

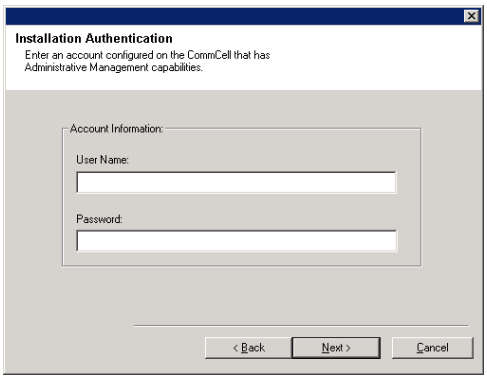

10. Enter the following:

- 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.

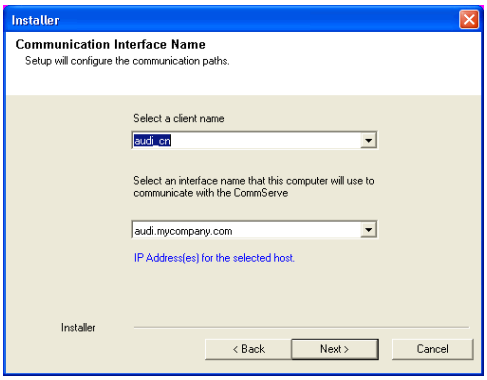

11. Select **Add programs to the Windows Firewall Exclusion List**, if you wish to add CommCell programs and services to the Windows Firewall Exclusion List.

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.

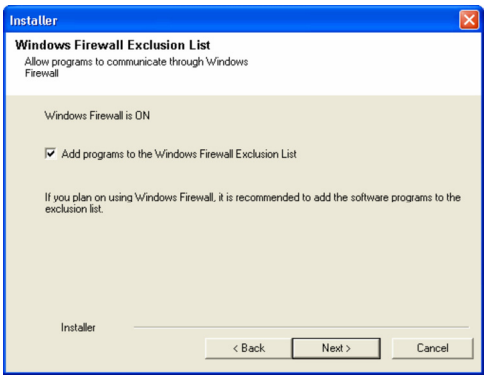

12. Specify the location where you want to install the software.

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.

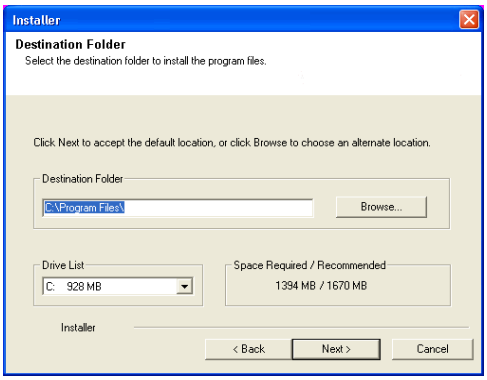

SCHEDULE AUTOMATIC UPDATE

13. 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.

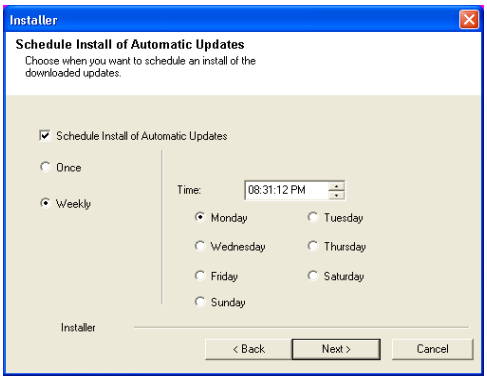

REPLICATION LOGS LOCATION

14. Specify a path on a local volume for the Replication Logs directory.

Click **Browse** to change directories.

Click **Next** to continue.

NOTES

- For CDR on a cluster, Replication Logs must be located on a local volume, not a volume which is part of the cluster resource group.
- There are several additional considerations for log file space and location; refer to Replication Logs when deciding on a suitable location for Replication Logs.

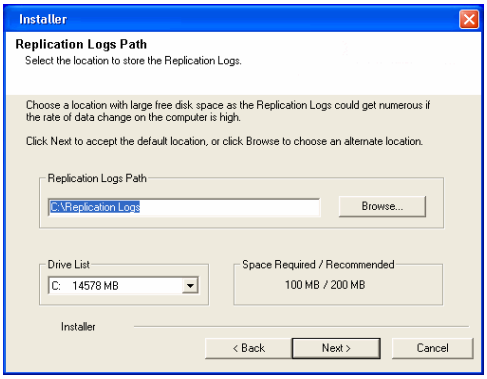

15. Setup reminds you that the computer must be restarted, after the installation completes, before you can use this Agent.

Click **OK** to continue.

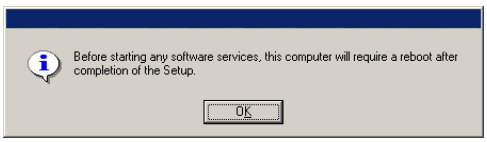

VERIFY SUMMARY OF INSTALL OPTIONS

16. 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.

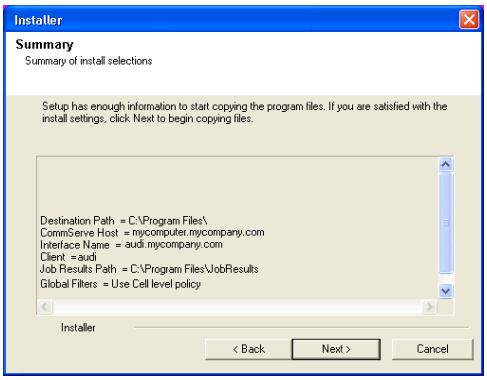

INSTALL REMAINING CLUSTER NODES

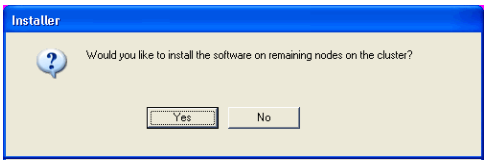

17. If you are installing/upgrading the software on the physical node in a clustered environment, use this option to install/upgrade the software on the remaining physical nodes of the cluster.

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

See Install/Upgrade Remaining Cluster Nodes for step-by-step instructions.

SETUP COMPLETE

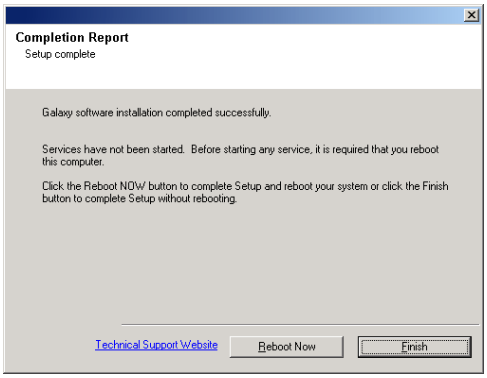

18. Select from the following:

- 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

- Perform this installation procedure for each physical node in the cluster, and then install the software in the cluster group; see Install ContinuousDataReplicator - Clustered Environment.
- Do not apply Updates or Service Packs until you have installed the Driver for ContinuousDataReplicator on all the physical nodes, and ContinuousDataReplicator on the cluster group.

Install ContinuousDataReplicator - Clustered Environment

TABLE OF CONTENTS

INSTALL REQUIREMENTS

The following procedure describes the steps involved in installing ContinuousDataReplicator in a clustered environment. Before you can install ContinuousDataReplicator in the cluster group in a clustered environment, you must first install the Driver for ContinuousDataReplicator on all physical nodes of the cluster. For step-by-step instructions, see Install ContinuousDataReplicator - Clustered Environment - Physical Node.

ContinuousDataReplicator 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 the software is installed in the cluster group, or you can choose to install any passive node(s) separately.

For an overview of deploying the software components in a clustered environment, see Clustering Support.

ContinuousDataReplicator is installed on at least two machines; the server from which you will replicate data (source computer) and the computer *to* which you will replicate the data (destination computer.) You may choose to perform additional installations, based on the Supported Configuration used in your particular environment.

Verify that the computers in which you are installing the software satisfy the minimum requirements specified in System Requirements - ContinuousDataReplicator.

Review the following Install Requirements before installing the software:

GENERAL

- If any of the computers in which you are installing this software have multiple Network Interface Cards (NIC) you must configure them so that the source and destination computers can communicate for replication activities. For more information, see Data Interface Pairs.
- Agents should only be installed after the CommServe and at least one MediaAgent have already 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.
- This version of the software is intended to be installed in a CommCell where the CommServe and MediaAgent(s) version is 9.0.0.
- 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 for the Agent.
- Verify that the software installation disc is appropriate to the operating system of the computer on which the software is being installed.
  Make sure that you have the latest software installation disc before you start to install the software. If you are not sure, contact your software provider.

CLUSTER

- 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

- 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 remotely installing to other physical nodes of the cluster, ensure you mount the software installation disc on a UNC path or a shared location (e.g., **`\\machine_name\shared_directory_name`**).

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.

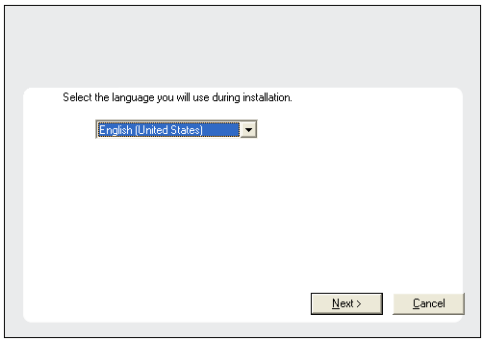

3. Select the option to install software on this computer.

   NOTES

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

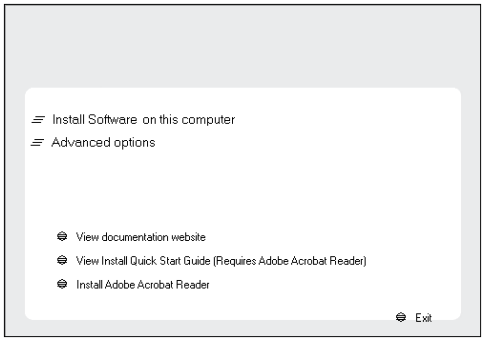

4. Read the license agreement, then select **I accept the terms in the license agreement**.

   Click **Next** to continue.

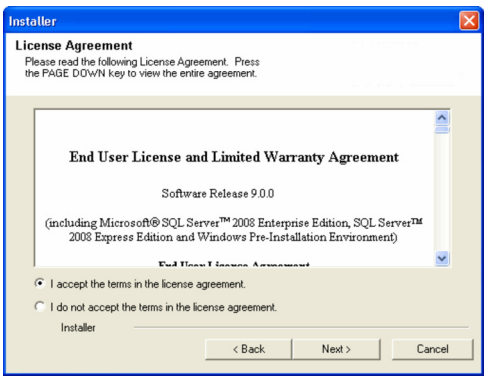

CLUSTER SELECTION

5. Select **Configure a Virtual Server**.

   Click **Next** to continue.

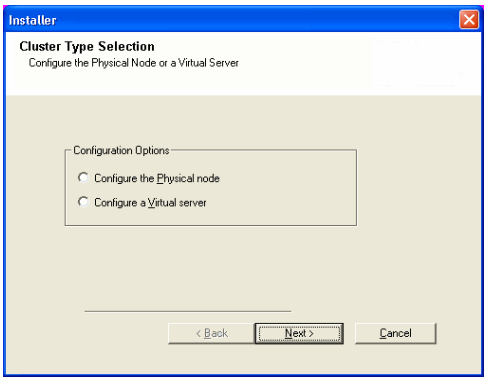

6. Select the disk group in which the cluster group resides.

   Click **Next** to continue.

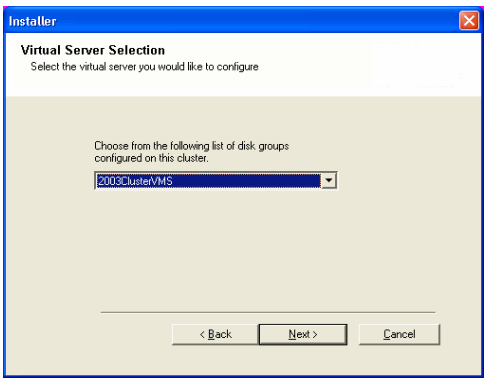

SELECT COMPONENTS FOR INSTALLATION

7. 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 ContinuousDataReplicator, expand the `Client Modules` folder and the `ContinuousDataReplicator` folder and select the following:

   - `ContinuousDataReplicator`

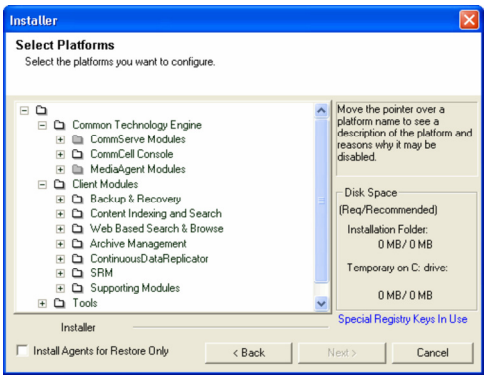

CONFIGURATION OF OTHER INSTALLATION OPTIONS

8. 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.

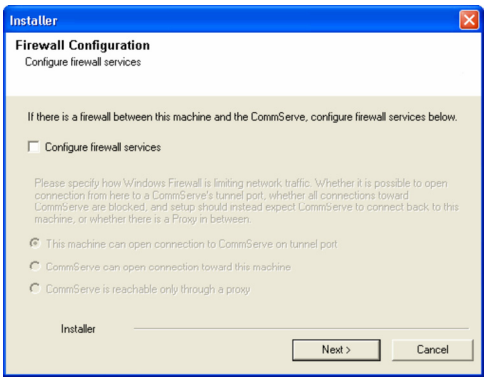

9. Enter the fully qualified domain name of the CommServe Host Name. This should be TCP/IP network name. e.g., computer.company.com.

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.

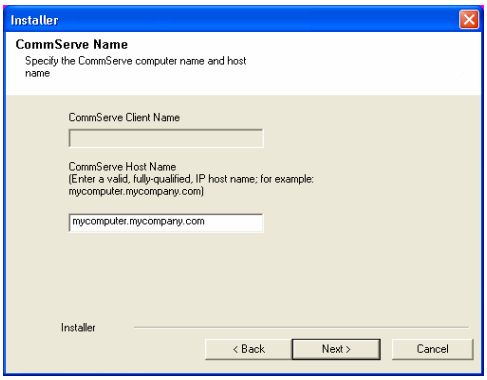

10. Enter the following:
    - The local (NetBIOS) name of the client computer.
    - The TCP/IP IP host name of the NIC that the client computer must use to communicate with the CommServe Server.

    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.

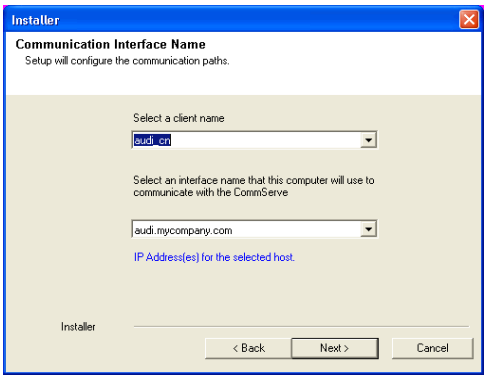

11. Select **Add programs to the Windows Firewall Exclusion List**, if you wish to add CommCell programs and services to the Windows Firewall Exclusion List.

    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.

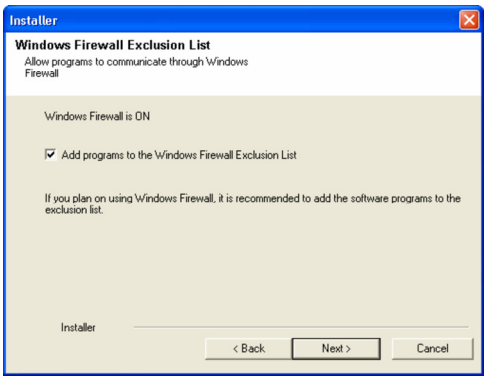

12. Enter the **username** and **password** associated with an external domain user account or a CommCell user account to authorize the installation of this agent.

    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.

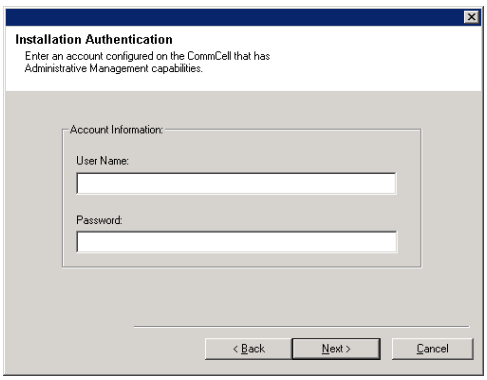

SCHEDULE AUTOMATIC UPDATE

13. 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.

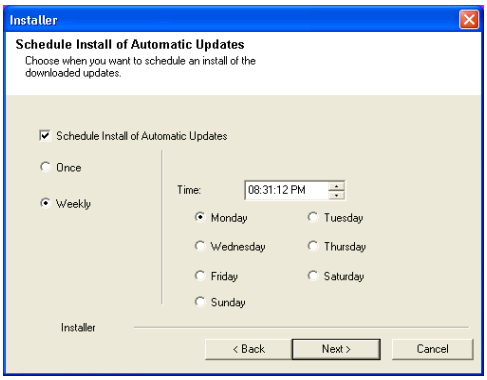

VERIFY SUMMARY OF INSTALL OPTIONS

14. 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.

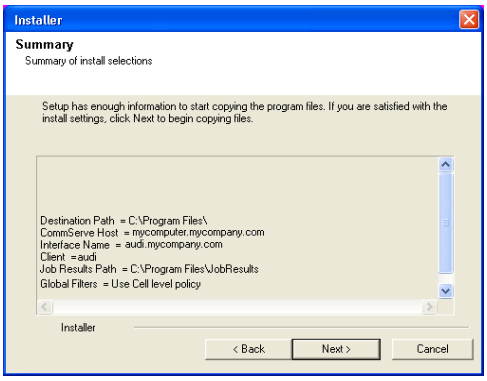

INSTALL REMAINING CLUSTER NODES

15. To install/upgrade the software on the remaining nodes of the cluster, click **Yes**.

To complete the install for this node only, click **No**.

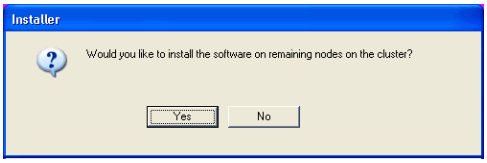

16. Select cluster nodes from the **Preferred Nodes** list and click the arrow button to move them to the **Selected Nodes** list.

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.

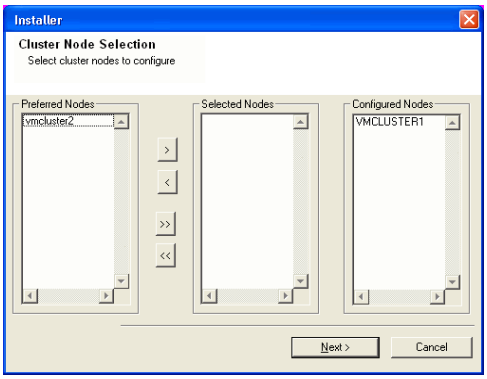

17. Type the **User Name** and **Password** for the Domain Administrator account, so that the installer can perform the remote install/upgrade of the cluster nodes you selected in the previous step.

Click **Next** to continue.

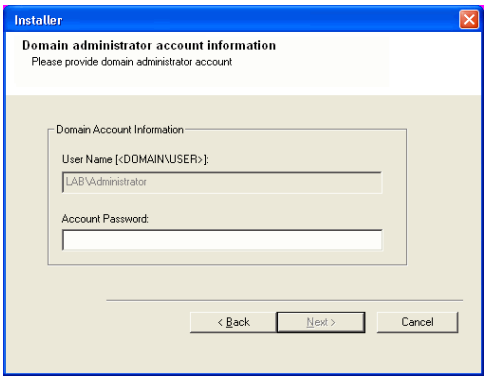

18. The progress of the remote install for the cluster nodes is displayed; the install can be interrupted if necessary.

    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.

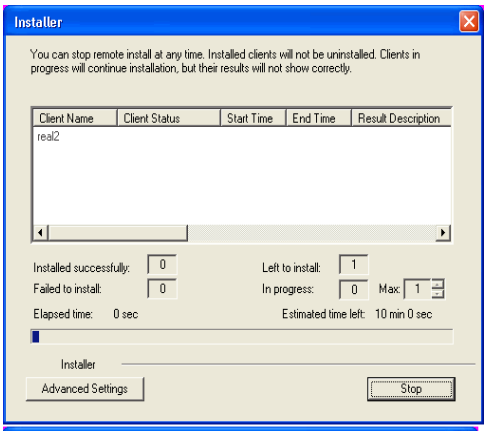

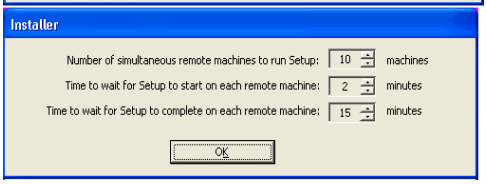

19. Read the summary for remote installation to verify that all selected nodes were installed successfully.

    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.

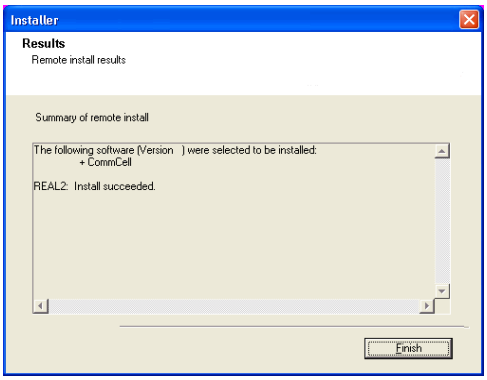

SETUP COMPLETE

20. Setup displays the successfully installed components.

    NOTES

    - The **Setup Complete** message displayed on your screen will reflect the components you installed, 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.
    - If **Reboot Now** button is displayed make sure to reboot the computer before performing any other operations from the computer.

    Click **Finish** to close the install program.

    The installation is now complete.

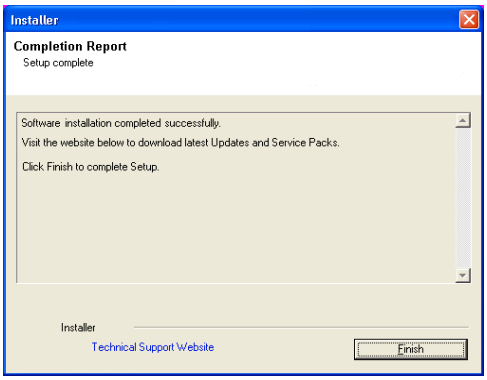

21. Note that the following will automatically be filtered:
    - The physical cluster nodes will automatically filter all shared disk resources.
    - Cluster group will automatically filter all physical disk resources.

    NOTES

    - For information on adding your own filters, see Add Filters for a Replication Set.

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 ContinuousDataReplicator - UNIX

TABLE OF CONTENTS

WHERE TO INSTALL

Install the software directly on the computer hosting the global zone or the Unix Server that you wish to protect. Make sure the computer satisfies the minimum requirements specified in the System Requirements.

On the source computer, the ContinuousDataReplicator can be used to replicate data from any non-global zone.

INSTALL REQUIREMENTS

The following procedure describes the steps involved in installing ContinuousDataReplicator. ContinuousDataReplicator is installed on at least two machines; the server *from* which you will replicate data (source computer) and the computer *to* which you will replicate the data (destination computer.) You may choose to perform additional installations, based on the Supported Configuration used in your particular environment.

Verify that the computers in which you are installing the software satisfy the minimum requirements specified in System Requirements - ContinuousDataReplicator.

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 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:
```

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.

```
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

3) Exit this menu

Your choice: [1]
```

5. 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 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]
```

6. 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**.

```
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]
```

7. 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.

```
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]
```

8. 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.

```
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

9. Enter the number corresponding to the **CVGxCDR** 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

10. If you wish to install the agent software for restore only, enter **Yes** and press **Enter**

```
Do you want to use the agents for restore only without
```

to continue. See Installing Restore Only Agents for more information.

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

```
consuming licenses? [no]
```

11. 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 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.

```
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.
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
```

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.

```
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 ...
```

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.

```
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 ...
```

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.

```
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 ...
```

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.

```
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]
```

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.

```
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 ...
```

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**.

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.

```
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.

Please enter the port numbers.

Port Number for CVD : [8600]

Port Number for EvMgrC: [8602]
```

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.

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

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.

```
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:
```

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.

```
Enter your CommCell user name and password:

User Name :

Password :

Press <ENTER> to continue ...
```

CDR CONFIGURATION

21. Specify whether the computer on which you are installing CDR will be a source, a destination, or both.

```
Any machine participating in the CDR network, can be
either the source, the destination or both.

More work has to be done when installing CDR on the source
machines: we will have to configure a special file system
driver and select a dedicated directory where we will keep
the cache of the file system changes.

Please specify how Calypso CDR is going to be used on this
machine. Whether it will be a:

1) Source Machine
2) Destination Machine
3) Both Source and Destination Machine

Your choice: [3]
```

22. Specify a path for the Replication Logs directory.

NOTES

- There are several considerations for log file space and location; refer to Replication Logs when deciding on a suitable location for Replication Logs.

```
Since this machine is going to be used as the Replication
Source, Calypso CDR needs a dedicated directory where it
will maintain the cache of file system changes. The amount
of free space in that directory should be at least 2GB
(you will be able to customize it later).

CDR Cache Directory:
```

23. Specify, in MB, how much log file space should be allocated.

```
Please specify how much space Calypso CDR is allowed to
use in /test/cache for the caching purposes. Note that
Calypso will use up all of the allocated space fairly
quickly. Please don't enter more than you can really
spare.

You have 4398MB free in /test/cache.

CDR Cache Size, MB:
```

24. The install program now starts copying the software to the computer. The progress of the operation is displayed. If you are installing on a Linux computer and are prompted that the version of the kernel is not supported for the driver, boot one of the

```
.....
.....
.....
```

identified kernels and then press **Enter**.

```
.....

Successfully copied xx files

.....
.....
.....

Successfully installed CVGxCDR
```

SETUP COMPLETE

25. 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]
```

26. If this is the last package that you wish to install, enter the number corresponding to the **Exit** option and then press **Enter** to continue.

```
Calypso now supports the concept of instances, each of
which can be installed to a separate CommServe and be
totally independent of others. Instances are not the same
as the virtual machines on a cluster, in fact, you will be
given a choice to configure one or more virtual machines
later. Contrary to previous releases of Calypso (6.1 and
older), all virtual machines installed on the same host
now share the same set of binaries and services and one
situation where it is useful is when you want to back up
the same host to several CommServes.

You already have one instance configured.

1) Add agents to existing Instance001 going to CS
mycompany.company.com

2) Create a new instance of Calypso

3) Exit this menu
```

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.
- On the destination computer, data can be replicated to a non-global zone through a exporting device from the global zone. See Data Replication on Non-global Zones for more information.

Data Replication

Topics | How To

TABLE OF CONTENTS

OVERVIEW

Data Replication is the process of copying specified, file-level content from one computer, the source computer, to another, the destination computer. This is achieved through an initial transfer of the specified data, after which the replicated copy is kept updated in real time with any changes that are made to the data on the source computer. This replicated copy on the destination computer provides on-going, nearly-real-time disaster recovery protection for the source computer, unlike most data protection solutions which require significant time to perform a complete data protection operation. In addition, data replication provides a basis for additional data protection activities, such as Recovery Points (snapshots) and backups of Recovery Points.

The content for replication can be defined at the directory or volume level on a source computer and replicated to a destination computer. Once the initial transfer is complete, a driver on the source computer performs the following:

- continuously monitors changes to the files contained in the defined directories or volumes
- logs all new files, and changes to existing files
- automatically transfers the log to the destination computer, thus replicating all new files and changes to existing files, from the source computer to the destination computer in nearly real time. See Replication Logs for specific information about frequency and timing of data replication.

A persistent connection is used as a data transfer mechanism, optionally compressing and encrypting data across the network, and through this facility, the destination computer is kept in sync with the defined content on the source computer. If the connection is interrupted at any point, the log continues to be maintained on the source computer, and once the connection is restored, CDR will automatically re-sync with the destination computer, bringing the replica up-to-date. Note that re-syncing is time and disk space intensive, and thus to be avoided if possible. For some additional discussion of this subject, see Interruptions and Restarts. If multiple Replication Pairs are active, CDR uses multiple threads to perform these operations on all Replication Pairs in parallel. CDR operations on a T1 link are fully certified. The success of CDR operations on a slower link is not guaranteed.

SUPPORTED CONFIGURATIONS

Some of the scenarios for data replication are listed below, but this is not a complete list of all the possible data replication configurations.

- Direct Replication
- Fan-In Configuration
- Fan-Out Configuration

DATA SYNCHRONIZATION

The following options can be used to perform data transfer from source to destination:

FULL RESYNC

Full Resync should be necessary only in cases when no data presented on destination. Full Resync copies all the files from the source to the destination computer. When you start Full Resync at Replication Set or Replication Pair level, you can specify Full Resync, causing the Replication Pair to begin at the Baseline Scan phase.

SMART SYNC

Smart Re-Sync is the default behavior of CDR when activities are interrupted and cannot be seamlessly restarted at the same point again. In this case all new/modified data will be transferred from source to destination.

OPTIMIZED SYNC

If replication is interrupted and there is a chance that the data on the destination is manually partially deleted or modified etc., the destination path is considered as inconsistent and optimized sync is recommended to rebuild it again based on the current data in the source path with consideration of data which already presented on destination.

Optimized Sync is used to transfer the modified/new files on the source computer to the destination computer along with data missing on destination. In previous attempts of sync had failures these failures will be re-tried during running Optimized Sync.

Optimized can be used in the following scenarios

- If after interruption in replication the filtering option are modified, such as removing filter that was previously applied, pre-existing files become eligible for transfer to the destination
- If some data was partially modified on destination
- If previous sync had failures.

See Start Data Replication Activity for step-by-step instructions.

To perform the Optimize Sync, see Add a Replication Pair for step-by-step instructions.

To change the state of one or more Replication Pairs at once from the Replication Set level, see Change the state of Replication Pair for step-by-step instructions.

TWO-WAY SYNC

You can use CDR to perform a two-way sync of data between the source and destination computers. The two-way sync up ensures that the data on the source and destination client computers is same and in sync with each other. The advantage of this feature is that the changes made to the data on the destination computer are replicated to the source computer. See Two-Way Sync, for more information.

REPLICATION PREDICTION

Replication Prediction can be used to track the size of the data that has been added or modified for the time during which a pair is active and monitoring; for Windows file systems, monitoring is performed at the volume or folder level; for UNIX, monitoring is performed at the file system level. This information is used to estimate the amount of data throughput required per hour, day, etc., and thus whether the bandwidth of the current connection will be sufficient for the predicted data replication activity. For instance, to see how much data will be replicated for an Exchange Server during each workday or for the whole week, you can start monitoring all folders used by the Exchange Server (stores, logs etc.) After 24 hours or a week, you can check the size of data modified, and use that information to estimate bandwidth requirements.

Replication Prediction reports the following for each monitored folder, volume, or file system:

- the monitoring interval -- start and end time
- the size of the data changed, in bytes and MB

For step-by-step instructions, see Perform Replication Prediction.

REPLICATION LOGS

CDR maintain logs on the computer, logging all file write activity (new files and changes to existing files) involving the directories and volumes specified in the source paths of all the Replication Pair(s) on that computer. These replication logs are transferred to the destination computer and replayed, ensuring that the destination remains a real-time replica of the source. For more information, see Replication Logs.

THROTTLING

Throttling enables you to monitor and control the data replication activities. It also allows you to configure the rate of data transfer over the network, based on the throttling parameters. The various throttling options (including throttling amount and rules) can be configured. For more information, see Throttling.

ORPHAN FILES

Files that are in the destination directory, but not the source directory, are orphan files. You can choose to ignore, log, or delete such files that are identified in the destination path; these settings are configured in the Orphan Files tab of the Replication Set Properties.

To configure Orphan File settings, see Configure Orphan File Processing for step-by-step instructions.

To view Orphan Files, see View Orphan Files for step-by-step instructions.

THINGS TO CONSIDER

- A file that is created on the source and is then deleted before it has been replicated, will still be created on the destination and then deleted. This is because both the creation and deletion of the file are captured in the log file, and this will be replayed on the destination computer. These are not treated as Orphan Files.
- A renamed file will be replicated to the destination as a new file. The previous copy with the old name will remain on the destination and be treated according to your Orphan Files settings.
- If you change the orphan file settings for an existing Replication Set, the change will only affect Replication Pairs that are created after the change, or Replication Pairs that are aborted and restarted. Currently active Replication Pairs will not be affected by the change until they are aborted and restarted.
- It is strongly recommended that you do not replicate to the root of the destination filer or the filer volume. If for any reason you need to replicate to the root of the volume then ensure that the Orphan File Processing is turned off from the Replication Set Properties.

DATA REPLICATION MONITOR

Replication is a continuous activity and details of on-going data replication activity is shown in the Data Replication Monitor in the CommCell Console. The process of starting data replication with CDR involves several job phases, as follows:

- Baselining
- SmartSync
- Replication

For more detailed information about Job Phases and Job States, see Monitoring Data Replication.

All other job-based activity, such as Recovery Point creation, is reflected in the Job Controller. See Controlling Jobs in Job Management for comprehensive information.

OUT OF BAND SYNC

In cases where large amounts of data must be transferred from the Source computer to the Destination computer during Baselining, but the connection between the source and the destination is constrained, such as a slow WAN connection, you may not want to begin replication using the Baselining Phases. You may prefer, for instance, to back up the source and restore it to the destination to effect the initial transfer of data.

To perform the initial transfer of data without using baseline, see Out Of Band Sync from the Replication Set for step-by-step instructions. After the transfer of data, start the Replication Pair with Start, so that only the data that is new or modified since the backup will need to be replicated.

REPLICATE THE DESTINATION DATA BACK TO THE SOURCE COMPUTER (WINDOWS ONLY)

It is recommended that you keep the following in mind when performing the replicate data back to the source computer:

- If data has been damaged on the source computer, perform a Copyback from the Live Copy on the destination, without **Overwrite existing data...** selected. See Copy Back File System Data from a Recovery Point or the Live Copy.
- In a case of failure of the source computer, the Replication Pair(s) can be aborted, and the data on the destination computer can be used as the primary data set. Once the problem is solved on the original source computer, the Replication Pair(s) can be created in reverse, replicating the new and modified data back to the source computer, using Smart Re-Sync.

  To limit the replication to only the data newly created or modified on the replica while it was being used as the production data set, you must save the current USN (Unique Sequence Number) on the destination volume(s) before actually using them as the production data set. This will ensure that when you start the Replication Pair later to replicate data back to the source computer, CDR can use Smart Re-Sync, beginning from the USN that was saved.

To Replicate the Destination Data Back to the Source Computer, see Replicate the Destination Data Back to the Source Computer for step-by-step instructions.

IMPORTANT CONSIDERATIONS

It is recommended that you keep the following in mind when performing data replication:

GENERAL

- Ensure that the destination volume has sufficient space for all the data that will be replicated to it. If you are replicating data from multiple source volumes to the same destination volume (Fan-In), ensure that the destination volume is sufficiently large for the data which will be replicated from all the source volumes. If you are creating Recovery Points, you must also account for the space requirements of the snapshots that will be created on the Destination; see Recovery Points - Snapshot space requirements.
- Individual failed files or folders will not necessarily fail the replication job. Such individual failures may just be logged and the data replication job will continue. Check the logs periodically for such failures. See View the Log Files of an Active Job. In some cases, the nature of such failures during replication may have an underlying cause which would in turn cause CDR to switch to SmartSync, or Abort replication altogether.

- In a case of failure of the source computer, the data on the destination computer can be used temporarily as the primary data set. Once the problem is solved on the original source computer, the new and modified data can be replicated from the destination computer back to the source computer. For more information, see Replicate the Destination Data Back to the Source Computer.
- If a SAN volume that is a source for any Replication Pair(s) is disconnected and re-connected again, you must abort and restart at least one of the Replication Pairs on the source computer.
- On both the source and destination computers, it is recommended that you Configure Throttling for CDR Replication Activities.
- It is recommended that you also Configure Alerts. For more information, see Application Management Alerts for CDR and Job Management Alerts for CDR.

WINDOWS

- When you replicate data that was encrypted on the source computer, it will not be accessible on the destination computer. To access the data, you must use Copyback to recover the data to the source computer, where you will be able to access it with the proper permissions. On the source computer, if you remove the encryption from the data after it has been replicated, the data will not be replicated again, so it will remain encrypted on the destination. Encrypted files are replicated in the Baseline and the SmartSync phases.
- The virtual memory paging file (`pagefile.sys`) must be configured on a local, fixed disk.
- When using QSnap, you may want to increase the minimum size of QSnap's COW cache beyond the default size, on both the source and destination computers, if sufficient space is available. Also, you may want to select an alternate location for the COW cache. For more information, see QSnap - Cache Considerations for ContinuousDataReplicator.
- When replicating application data, see Change Account for Accessing Application Servers.
- When using VSS or QSnap on a source computer it is recommended that you review Space Check and configure the `Disk Space Low` alert to provide warning that the source computer is running out of disk space, which will ultimately cause replication activity to be System Aborted.
- If Windows compression is set on root level of a driver letter, the compressed files will be replicated to destination as uncompressed files.

UNIX

- ACLs for AIX 5.3 cannot be replicated to a destination running AIX 5.2, as the ACL format is not backward compatible. However, ACLs from AIX 5.2 can be replicated to a destination running AIX 5.3.
- Sparse files attributes are not transferred during the Baselining and SmartSync phases; the files assume the attributes of regular files on the destination. During the Replicating phase, sparse files do retain their attributes on the destination.
- To use QSnap, before you can begin creating Replication Sets and Replication Pairs, you must first configure source and/or destination volumes as CXBF devices. For more information, see QSnap for ContinuousDataReplicator.
- To replicate files with non-ASCII character names, perform the procedure detailed in Configuring the Locale for Non-ASCII Characters.

CROSS PLATFORM REPLICATION

Cross Unix platform data replication is now supported. For example, you can replicate data from a AIX source computer to a Solaris destination computer, or Solaris to Linux, etc. However, ACLs and Extended Attributes will be lost.

REGISTRY KEYS FOR DATA REPLICATION

Use the following registry keys to modify the default behavior of the Data Replication:

TOPIC	REGISTRY KEY(S)	DESCRIPTION
Change Journal	`dwCJSizeAsPercentOfVolumeSize`	ContinuousDataReplicator on Windows and the Windows File System iDataAgent use Change Journal to track updates made to Windows File Systems. On very large or very busy file systems, it may be necessary to increase the size of the change journal in cases where the agent or enabler is performing full scans too frequently. You can control the amount of volume space that is allocated for Change Journal when it is created by using the `dwCJSizeAsPercentOfVolumeSize` registry key value.
Pipeline Buffer Size	`PipelineBufferSizeInKiloBytes`	Replication performance is used to increase the speed at which data is replicated. The pipeline buffer size can be reconfigured from the default size of 64KB up to a maximum of 256KB (in increments of 32KB) using the `PipelineBufferSizeInKiloBytes` registry key. See `PipelineBufferSizeInKiloBytes` in Registry Keys for more information.
Connection Attempt	`MaxConnectionAttempts`	When communication is interrupted between the source and destination computers, the source computer will make 30 attempts (this default number can be changed using the `MaxConnectionAttempts` registry key) to reconnect to the pipeline, after which the Replication Pair(s) will show a state of Failed. Each connection attempt takes several minutes, an interval which is neither programmatic nor configurable.
Access Control Files	`nDoNotReplicateACLs`	For Windows, the `nDoNotReplicateACLs` registry key can be used to disable the replication of the security stream of files. This stream includes user and group access control list (ACL) settings for file access. If this registry key is not present, ACLs will be replicated.

Back to top

Data Replication - How To

Topics | How To

TABLE OF CONTENTS

START DATA REPLICATION ACTIVITY

Before You Begin

- Review Data Replication.

Required Capability: Capabilities and Permitted Actions

To start Data Replication Activity:

- In the CommCell Browser, click the **ContinuousDataReplicator**.
- Select **Replication Set**, right-click a Replication Pair, and select any one from the following options:

  Start, or **Start Full Resync**.

CHANGE THE STATE OF REPLICATION PAIR

Before You Begin

- Review Data Replication

Required Capability: Capabilities and Permitted Actions

To change the state of a Replication Pair:

- You can change the state of one or more Replication Pairs at once from the Replication Set level.

1. In the CommCell Browser, right-click a Replication Set in the source machine, and select **All Tasks**, then select **Change State**.
2. In the **Change States** dialog box, select the Replication Pair(s), and click **Start**, **Start Full Resync, Suspend**, **Resume**, or **Abort**.
3. When your changes are complete, click **Close**.

CONFIGURE ORPHAN FILE PROCESSING

Before You Begin

- Review Data Replication - Orphan Files

Required Capability: Capabilities and Permitted Actions

To configure orphan file processing for a Replication Set:

1. From the CommCell Browser, right-click the Replication Set on a source computer, and select **Properties**.
2. From the Replication Set Properties (Orphan Files) tab, perform the following selections:
   - Select **Log orphan file names** to create a log of all orphan files found on the destination computer. If this option is not selected, orphan files will not be identified on the destination.
   - Select **Automatically deleted orphan files** to automatically log and delete all orphan files found on the destination.
3. Click **OK**.

VIEW ORPHAN FILES

Before You Begin

- Review Data Replication - Orphan Files

Required Capability: Capabilities and Permitted Actions

To view the orphan files for a Replication Pair:

1. In the CommCell Browser, right-click the ContinuousDataReplicator icon on a source computer, and from **All Tasks**, select **Data Replication Monitor**.
2. In the Data Replication Monitor, right-click a Replication Pair and select **View Orphan Files**.

PERFORM REPLICATION PREDICTION

Each listed command should be run from a command line prompt, in the folder where the base package resides.

Before You Begin

- Review Data Replication - Replication Prediction

Required Capability: Capabilities and Permitted Actions

START MONITORING

All active Replication Pairs are automatically monitored from the time they are configured. To monitor an object not configured using CDR, you will need to add the path of the object to the list of monitored objects by executing this command.

Windows: `predict -folder <full path and folder name>`

Examples:

- `predict -folder G:\data`
- `predict -folder G:\`

UNIX: `cdrp -start <fname>`

VIEW DATA FOR A MONITORED VOLUME, FOLDER, OR MOUNT POINT

View the data for a monitored volume, folder, mount point, or file system by executing this command.

Windows: `predict -getdata <full path and folder name>`

If you specify a volume for this command, all monitored paths in this volume will be reported.

Examples:

- `predict -getdata G:\data`
- `predict -getdata G:\`

Sample Output:

```
predict -getdata L:\

Vol Name: L:\
Folder \test1\
Monitored Interval From <time_stamp> to <time_stamp>
Monitored area Bytes Changed 00863744 Change in MB = 0.823730 MB

Folder \test\
Monitored Interval From <time_stamp> to <time_stamp>
Monitored area Bytes Changed 60018688 Change in MB = 57.238281 MB
```

UNIX: `cdrp -q[uery] [[-c] <fname>]`

Queries replication statistics for the file system that the specified file belongs to. If a file name is not provided, the program will enumerate all mounted watched file systems and will provide statistics for each of them. You may pass a "`-c`" option along with the file name. This will make the tool perform continuous queries -- refreshing the screen every several seconds.

RESET DATA FOR A VOLUME, FOLDER, OR MOUNT POINT

Reset the data for a monitored volume, folder, mount point, or file system by executing this command.

Windows: `predict -getdata <full path and folder name> -c`

If you specify a volume for this command, the data for all monitored paths in this volume will be reset.

Examples:

- `predict -getdata G:\data -c`
- `predict -getdata G:\ -c`

UNIX: `cdrp -r[eset] [<fname>]`

Resets replication statistics for the file system that the specified file belongs to. If a file name is not provided, the program will reset replication statistics for all mounted, monitored file systems.

STOP MONITORING

For objects not configured using CDR, stop monitoring by executing this command.

Windows: `predict -remove <full path and folder name>`

Examples:

- `predict -remove G:\data`
- `predict -remove G:\`

UNIX: `cdrp -stop <fname>`

OUT OF BAND SYNC FROM THE REPLICATION SET

Before You Begin

- All replication pairs within the replication set must be in new or stopped state.

Required Capability: Capabilities and Permitted Actions

To perform the Out Of Band Sync for a Replication Pair at Replication Set level:

1. In the CommCell Browser, expand the Client computers and locate the ContinuousDataReplicator source client.
2. Right-click the Replication Set, and from **All Tasks**, select **Out Of Band Sync**. **Out Of Band Sync** dialog box will be prompted.
3. Select **Setup for incremental replication without running a full sync** option to replicate the data without running a full sync.
4. Ensure that you copied the initial backup data on the destination computer, and then select **Enable Incremental Sync** option to start the Replication Pair with Smart Sync, so that only the data that is new or modified since the backup will need to be replicated.

OUT OF BAND SYNC FROM A REPLICATION PAIR

Before you Begin

- To perform Out of Band Sync, replication pair must be in new or stopped state.

Required Capability: Capabilities and Permitted Actions

To perform the Out Of Band Sync for a Replication Pair:

1. In the CommCell Browser, expand the Client computers and locate the ContinuousDataReplicator source client.
2. Right-click a Replication Pair, and select **Out Of Band Sync**. **Out Of Band Sync** dialog box will be prompted.
3. Select **Setup for incremental replication without running a full sync** option to replicate the data without running a full sync.
4. Ensure that you copied the initial backup data on the destination computer, and then select **Enable Incremental Sync** option to start the Replication Pair with Smart Sync, so that only the data that is new or modified since the backup will need to be replicated.

REPLICATE THE DESTINATION DATA BACK TO THE SOURCE COMPUTER (WINDOWS ONLY)

Required Capability: Capabilities and Permitted Actions

To temporarily use the replica as the production data set:

1. On the Destination computer, query and record the USN for each Replication Pair destination volume using the **`CDRXHelp.exe -readUSN -cn <ClientName> -vm <InstanceName*> -path L:\`** command, which returns the current USN of the specified volume. All Replication Pairs using this volume as a destination will share the same USN.

   Example:

   There are three Replication Pairs:

   ```
   E:\ => L:\E_drive
   F:\ => L:\F_drive
   N:\ => M:\N_drive
   ```

   Query the USNs of the `L:` and `M:` volumes (the first two Replication Pairs use the same destination drive, L:, and will thus use the same USN) by typing the following commands:

   - **`<software_installation_path>\Base> CDRXHelp.exe -readUSN -cn <ClientName> -vm <InstanceName*> -path L:`**

     Sample output:

     ```
     CDRXHelp 1.0 - CDR External Base Line Helper Tool
     ```

```
Path: L:\
USN: 0f4f5a68
```

- **<*software_installation_path*>\Base> CDRXHelp.exe -readUSN -cn <ClientName> -vm <InstanceName*> -path M:**

  Sample output:

  ```
  CDRXHelp 1.0 - CDR External Base Line Helper Tool
  Path: M:\
  USN: 0f4f6b46
  ```

  Record both USNs for use later in this procedure.

2. The replica on the Destination computer can now be used as the production data set.

When the original Source computer is operational again, proceed to the next section.

To replicate data back to the original Source computer:

1. On the Source computer, record your Replication Pair configuration information, then delete the Replication Pair(s) and Replication Set(s). See Delete a Replication Pair and Delete a Replication Set.
2. On the Destination computer, create a Replication Set and Replication Pair(s) -- reversing the source and destination for each Replication Pair that existed on the original Source computer -- and determine the Pair ID(s) using the **CDRXHelp.exe -cn <ClientName> -vm <InstanceName*> -pairs** command. See Create a Replication Set and Add or Edit a Replication Pair.

   Do not start the Replication Pair(s) yet.

   Example:
   Using the same example from Step 2, create three Replication Pairs on the Destination computer:

   ```
   L:\E_drive => E:\
   L:\F_drive => F:\
   M:\N_drive => N:\
   ```

   Determine the Replication Pair ID(s) using the **CDRXHelp.exe -cn <ClientName> -vm <InstanceName*> -pairs** command:

   <*software_installation_path*>\Base> CDRXHelp.exe -cn <ClientName> -vm <InstanceName*> -pairs

   Sample output:

   ```
   REPID PAIRID  SrcPath      DestPath
   91    177     L:\E_drive   E:\
   91    180     L:\F_drive   F:\
   92    183     M:\N_drive   N:\
   ```

3. On the Destination computer, set the USN for each Replication Pair to the one returned for the original Replication Pair's destination volume (the USNs you recorded in Step 2) by executing the following command:

   `CDRXHelp.exe -setUSN -cn <ClientName> -vm <InstanceName*> -pID <PAIRID> -USN <some_USN>` - sets a Replication Pair with a Pair ID `<PAIRID>` to a specific USN, `<some_USN>`.

   > If you miss this step your Replication Pair will have the default USN of `0` and when you start the Replication Pair, a Full Re-Sync will be performed.

   Example:

   Using the same example from Steps 2 and 4, execute the following three commands:

   <*software_installation_path*>\Base> CDRXHelp.exe -setUSN -cn <ClientName> -vm <InstanceName*> -pID 177 -USN 0f4f5a68
   <*software_installation_path*>\Base> CDRXHelp.exe -setUSN -cn <ClientName> -vm <InstanceName*> -pID 180 -USN 0f4f5a68
   <*software_installation_path*>\Base> CDRXHelp.exe -setUSN -cn <ClientName> -vm <InstanceName*> -pID 183 -USN 0f4f6b46

   Sample output:

   ```
   CDRXHelp 1.0 - CDR External Base Line Helper Tool
   PairId: 177
   Old JournalId: ffffffffffffffff
   Old USN: ffffffffffffffff
   New JournalId: 1c625d8d392c299
   New USN: 0f4f5a68
   ```

4. By default, orphan files will be automatically deleted. If there were files on the original Source computer which were never replicated to the Destination computer, you may not want such files to be deleted as orphans. See View Orphan Files and Data Replication - Orphan Files.
5. On the Destination computer, start the Replication Pair(s) using Smart Re-Sync, by using the **Start** command. See Start/Suspend/Resume/Abort Data Replication Activity.

6. Baselining will begin; all the data added, deleted, or modified since the saved USN will be copied back to the original Source computer.
7. When the Baselining phases end, stop using the Destination computer as a production server and create a Consistent Recovery Point to ensure all the data has been replicated back to the source. See Create a Recovery Point.
8. On the Destination computer, abort and delete the Replication Pair(s). See Start/Suspend/Resume/Abort Data Replication Activity and Delete a Replication Pair. Also, delete the Replication Sets; see Delete a Replication Set.

   When the Replication Set is deleted, any Recovery Points or Consistent Recovery Points that were created from the Destination will be deleted as well, which is necessary in order for SmartSync to succeed when you start replication on the original Source computer again. If you want to preserve the data in such Recovery Points, back them up before deleting the Replication Set. See Back up a Recovery Point.

9. On the original Source computer, recreate the original Replication Set and Replication Pair(s). See Create a Replication Set and Add or Edit a Replication Pair.

   Do not start the Replication Pair(s) yet.

10. To avoid transferring all the data again, set the USN for each Replication Pair before you start replication. Perform the following on the original Source computer:
    - Determine the Replication Pair ID(s) using the **`CDRXHelp.exe -cn <ClientName> -vm <InstanceName*> -pairs`** command as shown in Step 3. For this example, assume that the command returned Pair IDs of `200`, `205`, and `210`.
    - Set the USN for each Replication Pair to the current USN for its source volume, by executing the following commands:

```
<software_installation_path>\Base> CDRXHelp.exe -setUSN -cn <ClientName> -vm <InstanceName*> -pID 200 -USN 200
<software_installation_path>\Base> CDRXHelp.exe -setUSN -cn <ClientName> -vm <InstanceName*> -pID 205 -USN 205
<software_installation_path>\Base> CDRXHelp.exe -setUSN -cn <ClientName> -vm <InstanceName*> -pID 210 -USN 210
```

11. On the original Source computer, start the Replication Pair(s) using Smart Re-Sync, and begin using the Source computer as the production server again.

* InstanceName is the name used for a CDR instance (by default it is Instance001*).* If multiple instances of CDR are installed use the corresponding instance name, see Multi Instancing for more information.

Back to Top

Supported Configurations

Topics | How To | Related Topics

OVERVIEW

Some of the scenarios of data replication are listed below, but this is not a complete list of all the possible data replication configurations.

DIRECT REPLICATION

This is the most fundamental configuration for data replication. A single computer on the LAN or WAN has its data replicated to another computer, either local or remote. This provides protection of the source computer against catastrophic failure of the computer itself.

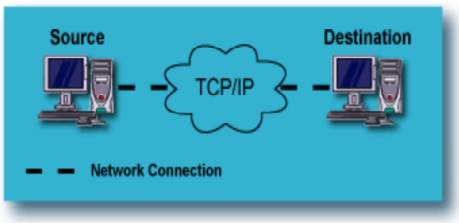

FAN-IN CONFIGURATION

In a Fan-In configuration, multiple computers on the LAN or WAN have their data replicated to a single computer, either local or remote. This provides protection of all of the source computers against catastrophic failure, while maximizing the use of resources by directing all the data to a single destination computer.

On Windows, most of the configuration of replication and Recovery Point options can be accomplished from the Fan-In tab of the Agent Properties on the *destination* computer, and these settings are automatically applied to all the source computer. On UNIX, replication and Recovery Point options must be configured on each source computer.

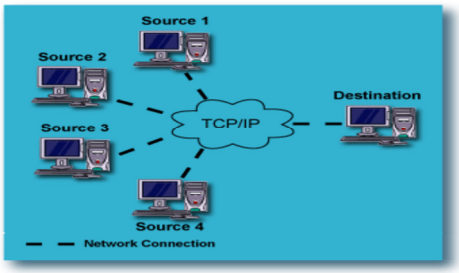

SCALABILITY

Although the scalability of a Fan-In setup can vary based on network and system resources, it is recommended that each Fan-In setup contains no more than 100 source clients.

For maximum performance and robustness, the total number of Replication Pairs configured for the same source volume should be kept to a minimum. If multiple Replication Pairs must be configured for the same source volume, the recommended upper limit is five.

In this configuration for data replication, see How to Setup Fan-In Configuration for more information.

RECOVERY POINTS FOR FAN-IN CONFIGURATIONS

Recovery Points created for a Fan-In configuration use VSS or ONTAP as the snap engine for creating snapshots. The use of snap engine is based on the destination being used. When the destination is a fixed volume then VSS is used and when the destination is a filer or iSCSI device then ONTAP is used for the creating snapshots. When destination volume is ONTAP LUN, Use ONTAP snapshot for ONTAP LUN destination option should be selected for creating snapshots with ONTAP snapshot engine. In case if option is not selected, VSS snapshot engine will be used.

Consider the following for ONTAP snapshots:

- Specify the user authentication details to be used for creating ONTAP snapshots. The user information must be specified both in the General tab of the replication set properties on the source and in the General tab of the agent properties on the destination.
- Manual mounting of ONTAP snapshots is not supported in a Fan-In configuration. If Auto-mount option is selected, ONTAP snapshots will be automatically shared.

HOW TO SETUP FAN-IN CONFIGURATION

The following section provides the steps required to use CDR for data replication and recovery in a Fan-In configuration based on multiple source computers and a single destination.

1. Install CDR software on all the designated source and destination computer.

   Verify that all selected computers meet the System Requirements. For more information see:

   - System Requirements - ContinuousDataReplicator
   - Deployment - ContinuousDataReplicator

   Also review the following:

   - Review Destination volume size in Data Replication - Important Considerations and Destination Computer Considerations.
   - The computer selected as the destination for a Fan-In configuration must have Windows 2003 or higher installed if you want to create Recovery Points.
   - Review Log Space Requirements, accounting for the number of source computers that will be sending Replication Logs to the single destination computer.

   To configure replicate application data, see Configure CDR to Replicate Application Data.

2. On Windows, configure replication and recovery point options. See, Configure Replication and Recovery Point.
3. Create a Replication Set.

   Note the following:

   - On UNIX, replication and Recovery Point options must be configured on each source Replication Set.
   - On Windows, the options specified in the previous step - configure replication and recovery point options - should not be specified when configuring the Replication Sets, e.g., Recovery Point options should only be set on the **Fan-In** tab of the destination **Agent Properties** dialog box, not on the **Replication Set Properties** dialog box on the source computer.

4. For each Replication Set, Add a Replication Pair.

   Ensure the following for each Replication Pair:

   - All source computers configured to Fan-In to the same destination computer.
   - Each Replication Pair must use the same destination volume but a different destination path.

5. Optionally, configure Throttling.

   Note the following:

   - Throttling can be used to allow the destination computer sufficient time to replay the logs it receives from all sources.

     On Windows, this can also help prevent the destination from running out of log space.

     For Fan-In configurations in particular, it is recommended that you configure the destination computer to start throttling the source computers when its log space runs low. For more information, see Replication Activity Throttling.

   - Configure the Throttling Amount on a source computer to limit the maximum amount of data (in MB/second) it can transfer to the destination, thus preventing a particular source from overwhelming the destination during periods of high I/O.

     However, if the source computer is throttled too much, it may not be able to transfer all its Replication Logs quickly enough, and it could run out of space. For more information, see Network Bandwidth Throttling.

6. Start Data Replication Activity.

HOW TO SETUP RECOVERY POINTS FOR A FAN-IN CONFIGURATION

WINDOWS

1. Recovery Points are configured on the destination. These settings will be applied to all Replication Sets for all source computers that use this computer as a destination. See Create a Recovery Point in a Fan-In Configuration for step-by-step instructions.

   Recovery Points created for a Fan-In configuration use VSS or ONTAP as the snap engine for creating snapshots. The use of snap engine is based on the destination being used. When the destination is a fixed volume then VSS is used and when the destination is a filer or iSCSI device then ONTAP is used for the creating snapshots. When destination volume is ONTAP LUN, Use ONTAP snapshot for ONTAP LUN destination option should be selected for creating snapshots with ONTAP snapshot engine. In case if option is not selected, VSS snapshot engine will be used.

Consider the following for ONTAP snapshots:

- Specify the user authentication details to be used for creating ONTAP snapshots. The user information must be specified both in the General tab of the replication set properties on the source and in the General tab of the agent properties on the destination.
- Manual mounting of ONTAP snapshots is not supported in a Fan-In configuration. If Auto-mount option is selected, ONTAP snapshots will be automatically shared.

2. Mount or Unmount the snapshots of a Fan-In Recovery Point, see Mount a Fan-In Recovery Point and Unmount a Fan-In Recovery Point. Note that the snapshots of a recovery point created for a fan-in configuration can be mounted but not shared.

UNIX

Recovery Points are configured on each Replication Set on each source computer. Thus, any Replication Set sharing the same destination can create Recovery Points. For instructions, see Configure Recovery Points and Create a Recovery Point.

FOR BOTH WINDOWS AND UNIX

- Optionally, you can Configure CDR for Backups of Recovery Points in a Fan-In Configuration.

  To perform backups of Recovery Points, you must also install the appropriate file system *i*DataAgent on the destination computer. Note that you cannot replicate Windows data to a UNIX computer, nor the converse. For the file system *i*DataAgent installation instructions, see:

  - Deployment - Windows File System *i*DataAgent
  - Deployment - AIX File System *i*DataAgent
  - Deployment - FreeBSD File System *i*DataAgent
  - Deployment - HP-UX File System *i*DataAgent
  - Deployment - Linux File System *i*DataAgent
  - Deployment - Solaris File System *i*DataAgent
  - Deployment - Tru64 File System *i*DataAgent
- View the Recovery Point Creation History.
- Run and review the CDR Copyback Job Summary Report.

RECOVERING DATA FOR A FAN-IN CONFIGURATION

Recovering data from Recovery Points or backups of Recovery Points is similar to Recovery Replicated Data. No additional or special steps are required. For more information on recovering data, see Recover Replicated Data. For Copyback procedures, see Copy Back File System Data from a Fan-In Recovery Point.

View the Recovery Point Copyback History.

Run and review the CDR Recovery Job Summary Report.

FAN-OUT CONFIGURATION

This configuration for data replication, referred to as "Fan-Out", adds significantly to the protection afforded to the source computer, because of the redundancy. A single computer on the LAN or WAN has its data replicated to multiple computers, any of which can be either local or remote. This provides protection against catastrophic failure of an entire site, as well as the source computer itself.

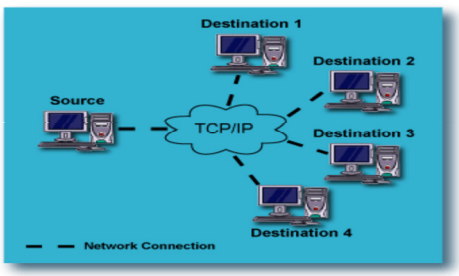

FAN-OUT CONSIDERATIONS

Consider the following for Fan-Out configuration:

- A snapshot of the source volume for each Replication Pair is created during the SmartSync Scan phase, which lead to significant space requirements in a Fan-Out configuration, since a separate Replication Set is required for each different destination. As a simple illustration, if you have 5 destinations for the same source path, thus 5 Replication Sets each having one Replication Pair, 5 snapshots will be created of the source volume. Further, if you have 10 volumes which are each being replicated to those 5 destinations, and you start replication for all of them simultaneously, 50 snapshots will be created during the SmartSync Scan phase. It is recommended in such circumstances to avoid starting all Replication Pairs simultaneously.
- The VSS cache can be configured using the **`vssadmin add shadowstorage`** command from a command line prompt. Refer to Microsoft documentation for details.
- QSnap COW Cache space considerations:
  - As the SmartSync phases complete for each Replication Pair, these snapshots are deleted by the system, but only in the order they were created for each volume. If the first snapshot of a given volume was created for a Replication Set using a slower destination computer, and it is the last one to complete the SmartSync phases, no other snapshots of that volume can be deleted until this first one is deleted.
  - By default, the location for the QSnap COW Cache is on the same volume of which the snapshot is being created, but you can also specify a separate volume to be used for all snapshots, in the Client Properties (Advanced) tab. See Change the COW Cache Location.
- For Windows, note that since each destination computers communicates with the source computer to indicate when it is finished with a log, a given log will not be automatically deleted on a source computer until *all* destination computers are finished with it. If one or more destinations are unavailable for any reason, (or planning to be) in Fan-Out scenarios for prolonged period of time, the Replication Pairs for that destination should be aborted, or the source will eventually run out of space as all the replication logs for the offline destination(s) accumulate.

Back to Top

Supported Configurations - How To

Topics | How To | Related Topics

CONFIGURE REPLICATION AND RECOVERY POINTS

Before you Begin

- Review Pre/Post Process
- Review Recovery Points

Required Capability: Capabilities and Permitted Actions

To configure replication and Recovery Point options on Windows:

1. From the CommCell Browser, right-click the ContinuousDataReplicator for the destination computer and select **Properties**.
2. In the **ContinuousDataReplicator Properties** dialog box, click the **Fan-In** tab.
3. Click the **Pre/Post Process** tab.
   - 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.

     Pre-Recovery Point Command.

     Post-Recovery Point Command.
   - In the User Account dialog box, select **Use Local System Account**, or select **Impersonate User** and enter a user name and password. Click **Ok**.
4. Click the **Recovery Points** tab.
   - **Enable the Fan-In** on destination computer.
   - Select Auto-mount snapshots after Recovery Point Creation option.
   - Select **Use ONTAP snapshot for ONTAP LUN destinations** option, to create ONTAP snapshots on iSCSI LUN.
   - Specify **Maximum Number of Recovery Points**.
5. Click **Storage Policy** tab.
   - Select **Backup Recovery Points** option.
   - Select the **Storage Policy**.
6. Click **Ok** to save your changes.

   Settings in these tabs will be applied to all Replication Sets for all source computers that use this computer as a destination.

Additional Recommendations

- It is recommended that you Configure Alerts. For more information, see Alerts.

  On a source computer it is recommended that you also see Space Check for the Quick Recovery and ContinuousDataReplicator Agents and configure the `Disk Space Low` alert to provide warning that the source computer is running out of disk space, which will ultimately cause replication activity to be System Aborted.
- Monitor Data Replication Activities.
- Run and review the CDR Replication Job Summary Report.

CREATE A RECOVERY POINT IN A FAN-IN CONFIGURATION

Before You Begin

- Review Using ContinuousDataReplicator in a Fan-In Configuration
- Recovery Points cannot be created until you Configure Recovery Points. Note that for Workstation Backup Agent, Recovery Point creation is enabled by default.

Required Capability: See Capabilities and Permitted Actions

To create a Recovery Point for a Fan-In configuration:

1. From the CommCell Browser, right-click **ContinuousDataReplicator** on the destination client, and select **Browse Destination Volumes**.
2. In the **Browse Destination Volumes** dialog box, right-click a destination volume, and select **Create Recovery Point**.
3. From the Create Recovery Point dialog box, select whether to create the Recovery Point immediately or schedule it.
4. In addition, you can select the option of backing up Recovery Points. If you select this option, click **Enter Backup Details** to select the type of backup as well as any Advanced Options.
5. When you are finished, click **OK**.
6. If you opted to schedule Recovery Point creation, the **Schedule Details** dialog box will appear. Specify the details for your schedule, click **Options** for additional choices. When you are satisfied, click **OK**.

DELETE A FAN-IN RECOVERY POINT

Before You Begin

- Review Using ContinuousDataReplicator in a Fan-In Configuration
- Deleting a Recovery Point will remove its corresponding snapshots as well as the Recovery Point Creation History.

Required Capability: See Capabilities and Permitted Actions

To delete a Fan-In Recovery Point:

1. From the CommCell Browser, right-click **ContinuousDataReplicator** on the destination client, and select **Browse Destination Volumes**.
2. In the **Browse Destination Volumes** dialog box, right-click a Recovery Point, and select **Delete**.

AUTOMATICALLY MOUNT SNAPSHOTS FOR FAN-IN RECOVERY POINT

Before you Begin

- Review How to Setup Fan-In Configuration
- You can configure CDR so that the snapshots that comprise the latest Fan-In Recovery Point will be automatically mounted when it is created. This will affect all Replication Sets that use the computer as a Fan-In destination. Alternately, once Recovery Points have been created, you can select an existing Recovery Point and mount its snapshots.

Required Capability: See Capabilities and Permitted Actions

To configure CDR to automatically mount snapshots for Fan-In Recovery Points:

1. From the CommCell Browser, right-click **ContinuousDataReplicator** on the destination client, and select **Properties**.
2. From the **Fan-In** tab, select the **Recovery Points** tab, and select the **Auto-Mount snapshots after Recovery Point Creation** option.
3. Click **Enter Mountpoints**.
4. In the Enter Mount Points dialog box, click the **Mount Point** column for each Destination Path and specify a drive letter (e.g., G:) to which the snapshot will be mounted.

   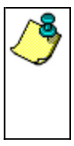

   - Ensure that you specify an *unused* drive letter.
   - Do not specify mount points for a UNC path. If you specify the mount point for a recovery point on a UNC path, this mount point is ignored when a recovery point is created on a UNC path.
   - When a Recovery Point snapshots created with ONTAP snap engine and **Auto-Mount snapshots after Recovery Point Creation** option is selected, share of snapshot will be created.

5. Click **OK**.

MANUALLY MOUNT OR SHARE THE SNAPSHOTS OF AN EXISTING FAN-IN RECOVERY POINT

Before you Begin

- Review How to Setup Fan-In Configuration

Required Capability: See Capabilities and Permitted Actions

To mount the snapshots of an existing Fan-In Recovery Point, on Windows:

1. From the CommCell Browser, right-click **ContinuousDataReplicator** on the destination client, and select **Browse Destination Volumes**.
2. In the **Browse Destination Volumes** dialog box, right-click a Recovery Point, and select **Mount**.
3. Specify a drive letter (e.g., `G:`) to which to mount the Recovery Point. Ensure that you specify an *unused* drive letter.
4. To share a snapshot on ONTAP, do the following, depending on which object you browsed:
   - In the **Browse** window, right-click the mounted snapshot to be shared, and select **Create Network Share**.
   - In the Enter Share Name dialog box, type a name for the share.
   - Select **Make share writeable** option. Click **OK**.
5. Close the **Browse** window.

UNMOUNT A FAN-IN RECOVERY POINT

Before You Begin

- Review Using ContinuousDataReplicator in a Fan-In Configuration.

Required Capability: See Capabilities and Permitted Actions

To unmount the snapshots of a Fan-In Recovery Point:

1. From the CommCell Browser, right-click **ContinuousDataReplicator** on the destination client, and select **Browse Destination Volumes**.
2. In the **Browse Destination Volumes** dialog box, right-click a Recovery Point, and select **Unmount**.
3. Close the **Browse Destination Volumes** dialog box.

CONFIGURE CDR FOR BACKUPS OF RECOVERY POINTS IN A FAN-IN CONFIGURATION

Before You Begin

- Review Using ContinuousDataReplicator in a Fan-In Configuration.

Required Capability: Capabilities and Permitted Actions

To configure CDR for backups of Recovery Points in a Fan-In configuration:

1. From the CommCell Browser, right-click **ContinuousDataReplicator** on the destination client, and select **Properties**.
2. From the Fan-In tab, select **Back up Recovery Points**, and specify a Storage Policy for the destination machine.
3. Click **OK** to save your changes.

- To perform backups of a Recovery Point, see Back up a Recovery Point.
- If you change the Storage Policy, the next backup will be converted to a full.

Back to Top

Replication Logs

Topics | How To | Related Topics

OVERVIEW

ContinuousDataReplicator (CDR) performs replication by logging all activities in the source computer and replaying the log in the destination. On the source computer CDR logs all file write activities (new files and changes to existing files) involving the directories and volumes specified in the source paths of all the Replication Pair(s). These replication logs are transferred to the destination computer and replayed, ensuring that the destination remains a *nearly* real-time replica of the source.

Note the following differences in behavior on Windows and on UNIX:

- On Windows, log files are transferred periodically, not continuously. The period is based on the following:
  - The amount of change activity. Logs are transferred when they reach 5MB in size.
  - At a specified time interval if there is not sufficient change activity to fill the log. This time interval is by default 15 minutes. If necessary, you can adjust this change interval. For step-by-step instructions, see Specify the CDR Log File Update Interval.
- On UNIX, logs are sent to the destination computer in real time, and replayed from the destination computer's memory.

DESTINATION COMPUTER CONSIDERATIONS

On Windows, replication logs are replayed serially on the destination computer - not in parallel. Hence, if you have many Replication Pairs configured to use the same destination computer, the destination computer should be able to receive and replay the replication logs at the same rate at which they arrive.

In order to avoid a backlog of replication log files causing the allocated log space to diminish to the point that throttling on the source computer(s) ensure that the destination computer has sufficient resources. This includes:

- processing power and memory
- I/O capacity
- disk space allocated for replication logs

On UNIX, replication logs for different Replication Sets are replayed in parallel. This is done using multiple replay threads on the destination computer. Ensure that the destination computer has sufficient resources. This includes:

- processing power and memory
- I/O capacity

Note that the logs are replayed from the destination computer's memory on UNIX computers. Hence disk space is not needed to store log files on the destination.

LOCATION OF REPLICATION LOGS

The location of the log files is specified during the the installation of ContinuousDataReplicator Agent. If necessary, this can be changed subsequently from the CommCell Console. See Specify CDR Log File Location on Source and Destination Computers for step-by-step instructions.

Consider the following when you specify the location for the replication logs:

- Select a suitable volume for the source replication logs, which has sufficient space for the expected amount of log file activity and accumulation, for your environment.
- Do not specify a removable drive as the replication log location.
- Replication log cannot be located on a UNC path.
- On a cluster, replication logs must be located on a local volume, not a volume which is part of the cluster resource group.
- On UNIX, changing the log location on the source computer will cause the Bull Calypso Replication Service (CVRepSvc) service to be recycled, this will cause

all Replication Pairs to stop replication briefly and then resume.

- On UNIX, replication logs/cache has to be on a mount point different than the ones monitored by the ContinuousDataReplicator driver(source paths). Otherwise, the replication may not run properly.

LOG SPACE REQUIREMENTS

Sufficient log file space is required on the source computer, and the destination computer for Windows. If a source computer runs out of log space (Windows) or attempts to create new entries in a log file before the old entries have been transferred (UNIX), logging will stop and all logs will be deleted. Thus, to avoid an interruption and restart, it is important to have sufficient space allocated for logs. For minimum log space requirements, see System Requirements - ContinuousDataReplicator. These minimums should be considered as a recommended starting point. Allow more space than recommended if it is available.

Consider the following when allocating space for logs:

- Log file sizes will reflect the actual size of the files added or the extent of changes made to files in the source path.
- The existing size of the data in the source path and the expected rate of additions and file changes, for all the Replication Sets and Pairs that will be configured on a given computer. Larger amounts of data, and high rates of change typically result in greater amounts of log space being required on the source computer.
- Capacity of, or throttling limits imposed upon the network used for replication. If network capacity is low, log space requirements will increase on the source, as data is not transferred quickly enough.
- Potential network outages or loss of connectivity. During such times, logs will continue to accumulate on the source, and sufficient space must be available to accommodate these circumstances. In the case of a source computer configured with multiple destinations, loss of connectivity with any one destination computer will prevent the logs from being deleted on the source computer in a timely manner. For additional information, see Fan-Out Considerations below.
- For a computer that serves both as a source computer and a destination computer, log space must be sufficient to accommodate the requirements of both of the capacities in which it serves.
- For a computer configured as a destination for multiple source computers (Fan-In), allocated log file space should be matched to the aggregate needs of all of its source computers. (On Windows, this is a disk space requirement; On UNIX, this is a memory requirement.)
- Utilize the Space Check feature to configure a Disk Space Low Alert (and Space Check Interval, if appropriate) for the source log volume, so that you will be notified when free space is running too low; refer to Space Check for the Quick Recovery and ContinuousDataReplicator Agents. For step-by-step instructions, see Configure Alerts.
- On Windows, configure the free disk space threshold for the source log volume so that data replication will be aborted well before the free space on the source log volume becomes too low, which can cause unpredictable results.

  To avoid this, set the Low Watermark for the source log volume on the source computer to 10% or higher. See Configure Throttling for CDR Replication Activities. In the event this threshold is reached, a **Log File Volume Reached Low Watermark** alert and **Failed Replication - Application Management** alert will be generated. For more information about this alert, see Alerts - Application Management. You will have to make sufficient space available on the source log volume, and manually start the Replication Pair with Full Sync.

DELETION OF LOG FILES

Each log will continue to be saved on the source computer until all destination computers signal that they have received that log and have finished replaying it. After this confirmation, the log will be marked for deletion on the source and the system will periodically delete such logs.

- On UNIX, the system reuses log files in a rotating manner once the allocated log file space becomes full, so the logs will never be deleted.
- On Windows, logs on the destination computer are marked for deletion after they have been replayed, and the system will periodically delete these files as well.

If by accident you manually delete a log file on the source computer from the Operating System, obviously it cannot be transferred to the destination and replayed. This will result in the destination no longer being completely in sync with the source. To resync the source and destination:

- On Windows, it will be necessary to abort activity for all affected Replication Pairs and restart them again using **Start Full Re-Sync**. For instructions on aborting and restarting replication, see Start/Suspend/Resume/Abort Data Replication Activity.
- On UNIX, the Replication Pairs will automatically SmartSync before returning to Replication.

Back to Top

Replication Log - How To

Topics | How To | Related Topics

Specify CDR Log File Location on Source and Destination Computers

Specify the CDR Log File Update Interval

SPECIFY CDR LOG FILE LOCATION ON SOURCE AND DESTINATION COMPUTERS

Before You Begin

- On Windows, this procedure can be performed on both the source and destination computer(s); On UNIX, this procedure only applies to a source computer.

Required Capability: Capabilities and Permitted Actions

To specify a location for the CDR log files:

1. In the CommCell Browser, right-click the Client on either a source or destination computer, and select **Properties**.
2. In the **Advanced** tab of the Client Computer Properties screen, type or browse to path for **CDR Log File Location**.
3. Click **OK** to save your changes.

SPECIFY THE CDR LOG FILE UPDATE INTERVAL

Before You Begin

- Review Replication Log.
- The time interval specified here will be used by all of the ContinuousDataReplicator clients in the CommCell.
- Changes to the time interval will take effect on each client when a Replication Pair is started or when the Replication Service is cycled.

Required Capability: Capabilities and Permitted Actions

To specify the update interval for CDR log files:

1. In the CommCell Console, click the **Control Panel** icon, then double-click **Job Management**.
2. In the Job Updates tab of the Job Management dialog box, specify the number of minutes in the **State update interval for ContinuousDataReplicator** field.
3. Click **OK** to save your changes.

Back to Top

Throttling

Topics | How To | Related Topics

OVERVIEW

Throttling enables you to fine-tune the data replication activities for your environment. Data Replication activities can be throttled using the following:

- Frequency of log file transfers. See Log files for more information.
- Space allocation on the source and destination computer. See Replication Activity Throttling for more information.
- Available network bandwidth. See Network Bandwidth Throttling for more information.

REPLICATION ACTIVITY THROTTLING

Replication Activity Throttling is supported only on Windows.

The following can be configured on the Source computer:

- Interruption of replication activity, based on the percentage of allocated log space remaining on the source computer. (See System Aborted.)

The following can be configured on the Destination computer, and is recommended; it will impact all source computers that use this destination computer:

- Throttle the source computer, based on the percentage of allocated log space remaining on the destination computer. When throttling is imposed, it will reduce the maximum transfer rate specified as the Throttling Amount for the source computer by 50%. Since this throttle is based on the value specified as the Throttling Amount for the source computer, if you do not specify a value, no throttling will be imposed.
- Stop the source computer from sending logs, based on the percentage of allocated log space remaining on the destination computer.

NETWORK BANDWIDTH THROTTLING

The following can be configured on the Source computer:

THROTTLING AMOUNT

In the throttling amount, maximum network transfer rate, is measured in megabits per second (Mbps).

On Windows, the setting for Throttling Amount will apply to each pipeline engaged in replication activity, not the aggregate of all pipelines for all Replication Pairs on the computer. With respect to the number of pipelines that will be active at any one time, review the following based on an example of having a Throttling Amount setting of 10Mbps:

- Each Replication Pair requires a pipeline during Baselining or SmartSync; if you have 3 Replication Pairs in the SmartSync phases at the same time, you have 3 pipelines active, and thus a 30Mbps maximum network transfer rate for the source client (3 pipelines at 10Mbps each.) After Baselining and SmartSync have completed for each Replication Pair, only 1 pipeline will remain in use, and the maximum network transfer rate will be reduced to 10Mbps.
- Fan-Out configurations require a pipeline for each destination client; if you have 5 Replication Sets configured to replicate data to 5 clients, you have 5 pipelines active, and thus a 50Mbps maximum network transfer rate for the source client (5 pipelines at 10Mbps each.)
- Other than the cases listed, Baselining or SmartSync, and Fan-Out, all other replication activity will require 1 pipeline, and thus the Throttling Amount specified will apply to the aggregate of all activity on the source computer, comprised of the replication activities of all Replication Pairs.

On UNIX, the setting for Throttling Amount will apply to the aggregate of all pipelines for all Replication Pairs on the computer. With respect to the number of pipelines that will be active at any one time, review the following based on an example of having a Throttling Amount setting of 30Mbps:

- Each Replication Pair requires a pipeline during Baselining or SmartSync; if you have 3 Replication Pairs in the SmartSync phases at the same time, you have 3 pipelines active, which will equally divide the specified 30Mbps maximum network transfer rate for the source client, for a maximum rate of 10Mbps each. After Baselining and SmartSync have completed for each Replication Pair, only 1 pipeline will remain in use, and the maximum network transfer rate of 30Mbps will apply to that pipeline.
- Fan-Out configurations require a pipeline for each destination client; if you have 5 Replication Sets configured to replicate data to 5 clients, you have 5 pipelines active, which will equally divide the specified 30Mbps maximum network transfer rate for the source client, for a maximum rate of 6Mbps each.

If your destination computer is serving that function for multiple source computer, you may need to set this limit on each of the source computers, such that the destination computer has sufficient time for log replay from all source computers. Note that the Throttling Amount must be specified first, before you can create Bandwidth Throttling Rules (see next item.)

BANDWIDTH THROTTLING RULES

The scheduled network throttling, specifying a time range, and a percentage of maximum transfer rate, specified as the Throttling Amount, during that time range. Once again, the limit set here will apply to the aggregate of all activity on the source computer, comprised of the replication activities of all Replication Pairs. This setting allows you to select high and low (or no) activity periods in a manner similar to the 'operation window' concept used when scheduling the activities of job-based Agents.

CONSIDERATIONS

- When configuring throttling, you should consider what unintended affects throttling might have on operations. As one example, if you have a source computer that has significant file write activity, and you impose network bandwidth throttling which makes it impossible to transfer the logs quickly enough to the destination computer to keep pace with the rate of change on the source computer, log file space requirements would increase dramatically on the source computer. In such a case, provision must be made for sufficient log file space, based on the expected activity and throttling.
- An example of a beneficial use of throttling involves configurations where multiple source computers are all configured to use the same destination computer. In this case, you may want to impose throttling on the source computers to allow the destination computer enough time to keep pace with all the log files it is receiving, and ensure sufficient log space on the destination computer as well to accommodate all of the logs it will be receiving.
- You can configure Alerts to be generated when throttling is imposed, or when 80 percent or more of a volume's disk space is consumed, for all of the client computer's volumes. For more information, see Alerts and Monitoring.
- On Windows in a clustered environment, when a cluster node is the active node for more than one virtual server at the same time, throttling rules are applied equally to all of the virtual servers hosted by that physical node, using the highest numbers specified for any one of them. For example, consider an active node hosting three virtual servers simultaneously, with throttling configured as follows on each of the virtual servers, VS1, VS2, and VS3:

Throttling Parameter	VS1	VS2	VS3
Throttling based on percentage of free log space on destination:	30%	35%	**40%**
Stop replication based on percentage of free log space on destination:	**80%**	70%	60%
Abort source based on percentage of free log space on source:	75%	**80%**	70%
Network Bandwidth Throttling amount:	10Mbps	40Mbps	90Mbps

Since throttling for *all* Virtual Servers will be based on the highest number specified for any *one* of them, all three Virtual Servers would be subject to the throttling numbers shown in bold, not necessarily the numbers specified individually. If throttling is imposed based on the destination computer running low on log space, in this example, when free log space reaches 40% on any virtual server, the maximum transfer rate will be reduced by 50% on each of the virtual servers -- to 5Mbps on VS1, 20Mbps on VS2, and 45Mbps on VS3.

For step-by-step instructions, see Configure Throttling for CDR Replication Activities.

Back to Top

Throttling - How To

Topics | How To | Related Topics

CONFIGURE THROTTLING FOR CDR REPLICATION ACTIVITIES

Before you Begin

- Review Throttling.

Required Capability: Capabilities and Permitted Actions

To configure throttling for CDR replication activities:

1. In the CommCell Browser, right-click the ContinuousDataReplicator icon of the source machine, and select **Properties**.
2. In the Operational Parameters tab of the CDR Properties screen, specify any of the following:

   Destination Computer: (On Windows only)

   - Minimum percentage of free disk space for log files, below which the rate at which log files are sent from the source machine(s) is throttled. This throttling will reduce the maximum transfer rate specified in **Throttling Amount** by 50%. If you do not specify a value in **Throttling Amount**, no throttling will be imposed.
   - Minimum percentage of free disk space for log files, below which to stop the source machine(s) from sending more log files.

   Source Computer:

- On Windows, the minimum percentage of free disk space for log files, below which data replication is aborted.
- Throttling Amount - maximum network transfer rate in megabits per second (Mbps). The Utilization Percentage specified in the **Edit Throttling Rule** dialog will be a percentage of the number entered here.
- Bandwidth Throttling Rules - click **Add** to configure the following in the **Edit Throttling Rule** dialog:

  Days of Week

  Start time

  End time

  Utilization Percentage - based on the specified Throttling Amount (see previous item)

3. Click **OK** to save your changes.

Back to Top

Recovery Points

Topics | How To | Related Topics

Overview

Recovery Point is set of snapshots of the data to preserve a point-in-time on the Destination. 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.

For any supported configurations snapshot-based Recovery Points can be created and backed up. In addition, Recovery Points can be mounted, and for Windows can also be shared, to make it available to users on the network.

The illustration on the right provides a high-level look at how data flows when creating Recovery Points and backup of Recovery Points.

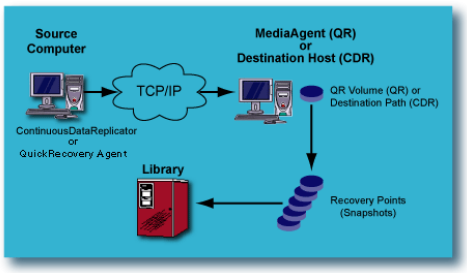

A Recovery Point is created at the Replication Set level. It will consist of a snapshot of every destination volume used by all the Replication Pairs that comprise the Replication Set.

Here is an example:

`repset1` is comprised of the following Replication Pairs:

- `reppair1`: destination is `F:\repset1_data\reppair1\`
- `reppair2`: destination is `F:\repset1_data\reppair2\`
- `reppair3`: destination is `G:\repset1_data\reppair3\`
- `reppair4`: destination is `H:\repset1_data\reppair4\`

Since the destinations used for all of the Replication Pairs in this Replication Set involve three volumes, namely `F:`, `G:`, and `H:`, on the destination computer, each Recovery Point for this Replication Set would contain three snapshots.

Recovery Point Types

Three types of Recovery Points can be used with ContinuousDataReplicator (CDR). They are:

Recovery Points

Recovery Points consist of snapshots created on the destination computer by the specified snapshot engine, without any reference to the state of the source computer. It simply represents a point-in-time on the destination computer, and is thus more useful for file system data than for application data.

CONSISTENT RECOVERY POINTS

A Consistent Recovery Point defines a point-in-time in which data is in a consistent state on the source computer. This ensures the data can be restored to that point-in-time. Hence it is more useful for application data than for file system data. Consistent Recovery Points are created as follows:

- On the source computer, the application server software is briefly quiesced and a marker is placed in the log file that is maintained on the source and transferred periodically to the destination computer.
- On the destination computer, when that marker is reached during replay of the log, a snapshot will be taken on the destination by the specified snapshot engine, which preserves a Consistent Recovery Point, and represents exactly the point-in-time on the source computer when the application was quiesced and the marker placed.

This process ensures the application can be restored to the exact point-in-time when the marker was placed on the source. Note that for non-integrated applications, CDR will not automatically quiesce the application server, but you can configure this behavior through the use of a quiesce and unquiesce script. When Consistent Recovery Point creation has been specified for a Replication Set, CDR will automatically check for the existence of quiesce/unquiesce batch files. If they are found, the script-based quiescing and unquiescing will be executed in conjunction with the placing of a marker in the Replication Log file on the source. Also, read Application Integration to understand how CDR replicates application.

You can create appropriate batch files using the name and location on the source computer given below. Note that you must use the Replication Set name to uniquely differentiate quiesce/unquiesce batch files from each other, since you may use such batch files for multiple Replication Sets.

- Quiesce batch file location and name on the source computer:

  <software installation path>\<Replication Set Name</>_Quiesce.bat

- Unquiesce batch file location and name on the source computer:

  <software installation path>\<Replication Set Name</>_Unquiesce.bat

RECOVERY POINTS FOR FAN-IN CONFIGURATION

Recovery Points created for a Fan-In configuration use VSS or ONTAP as the snap engine for creating snapshots. The use of snap engine is based on the destination being used. When the destination is a fixed volume then VSS is used and when the destination is a filer or iSCSI device then ONTAP is used for the creating snapshots. When destination volume is ONTAP LUN, Use ONTAP snapshot for ONTAP LUN destination option should be selected for creating snapshots with ONTAP snapshot engine. In case if option is not selected, VSS snapshot engine will be used.

Consider the following for ONTAP snapshots:

- Specify the user authentication details to be used for creating ONTAP snapshots. The user information must be specified both in the General tab of the replication set properties on the source and in the General tab of the agent properties on the destination.
- Manual mounting of ONTAP snapshots is not supported in a Fan-In configuration. If Auto-mount option is selected, ONTAP snapshots will be automatically shared.

PREREQUISITES

The following are the prerequisites for implementing Recovery Points.

INSTALL REQUIREMENTS

- The destination computer must meet the System Requirements for ContinuousDataReplicator. If you are using destination for Recovery Point creation, the destination computer must also meet the requirements for VSS or ONTAP.
- CDR software must be installed on the destination computer. If you are using QSnap for Recovery Point creation (Unix) or planning to create Consistent Recovery Points with Exchange Mining option enabled (Windows), QSnap software must also be installed on the destination computer as well.

  To perform backups of Recovery Points, the appropriate File System *i*DataAgent must be installed on both the source and destination computers.

  Refer to the following pages for more information on installing the software

  - Deployment - ContinuousDataReplicator
  - Deployment - QSnap
  - Deployment - Windows File System *i*DataAgent
  - Deployment - AIX File System *i*DataAgent
  - Deployment - FreeBSD File System *i*DataAgent
  - Deployment - HP-UX File System *i*DataAgent
  - Deployment - Linux File System *i*DataAgent
  - Deployment - Solaris File System *i*DataAgent
  - Deployment - Tru64 File System *i*DataAgent
- To use VSS writers for online quiesce of the Exchange Server or SQL Server on the source during the creation of Consistent Recovery Points, the **VSS**

Provider for CDR must be installed on the source computer.

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.

HOW TO SET UP RECOVERY POINTS

The following section provides the steps required to configure and use Recovery Points with CDR, based on a single source and single destination. If your environment uses a different scenario, adjust your steps accordingly.

Perform the following tasks to create Recovery Points:

1. If you wish to configure Consistent Recovery Points, you must first configure CDR to replicate data.

   See Configure CDR to Replicate Application Data for step-by-step instructions.

2. For both Recovery Points or Consistent Recovery Points, you must then configure the recovery points.

   See Configure Recovery Points for step-by-step instructions.

3. Next you must create the Recovery Points.

   For Direct Replication and Fan-Out Configurations, see Create a Recovery Point for step-by-step instructions.

   For Fan-in Configuration, see Create a Recovery Points in a Fan-in Configurations for step-by-step instructions.

You can also optionally perform the following operations:

- Pre/Post Commands for Recovery Points

  You can specify commands to run either before a Recovery Point is created and/or after a Recovery Point is created. For general information about Pre/Post commands, see Pre/Post Processes.

  See the following for step-by-step procedures for Pre/Post Processes with ContinuousDataReplicator:

  - Configure a Replication Set for Pre/Post Processes for Recovery Point Creation
  - Change Account for Executing Pre/Post Commands (Recovery Points)

- Mount or Share Recovery Point Snapshots

  The snapshots that comprise a Recovery Points can each be mounted/shared as a volume or mount point on the destination computer, and thus made available to users on the network.

  - On Windows, mounted volumes can also be shared on the network from the destination computer. These capabilities allow data to be copied from them by users with permissions to these volumes or shares, providing another means through which replicated data can be accessed easily.

    Mounting of Recovery Point snapshots can be configured to happen automatically whenever a Recovery Point is created, or you can select any existing Recovery Point and mount the snapshots on-demand. Snapshots can be mounted to a drive letter (e.g., `G:`) or a mount point (e.g., `G:\mountpoint1`) but note that the path for the mount point must exist already in order for the mount to be successful. In addition, you can also create a share on-demand from any mounted snapshot.

    When mounting a snapshot to a drive letter, you must ensure that it is an *unused* drive letter; attempting to mount a snapshot to a drive letter already in use by a local resources will fail. However, if the drive letter is already in use for a mapped network drive, the mount attempt will actually succeed, but produce unwanted results. As a result, for Windows 2003, the snapshot will be mounted, but not at the specified drive letter, and thus you will not be able to access it using Windows Explorer. The existing mapped network drive will continue to be accessible at that drive letter using Windows Explorer.

    The snapshots that comprise a Recovery Point can be shared, but they must be mounted first; conversely, they must be unshared before they can be unmounted. The snapshots of a Replication Set configured for Fan-in cannot be shared.

    To mount or share Recovery Points, see Mount or Share a Recovery Point for step-by-step instructions.

    To unmount or unshare a Recovery Points, see Unmount or Unshare a Recovery Point for step-by-step instructions.

    A Recovery Point cannot be deleted by the system if you have manually mounted or shared the snapshots that comprise it. For more information about why the system needs to delete Recovery Points, see Maximum Number of Recovery Points.

  - On UNIX, mount points are automatically created for all Recovery Point snapshots, and they are also automatically mounted; you have the option of creating a symbolic link to that mount path.

    On UNIX, CDR software automatically recognizes the File Systems configured on the destination computer and detects the appropriate snap engine. For more information on supported snap engines, see Snapshot Engines - Support.

MANAGE RECOVERY POINTS

All Recovery Points can either be scheduled or created on-demand. Consider the following sections for managing Recovery Points.

BACKUP RECOVERY POINTS

- Recovery Points can be configured to be automatically backed up when they are created by creating the necessary schedules. Also, you can select any existing Recovery Point and perform an immediate full backup on-demand. To perform a Backup of Recovery Points both the source and destination computers must have the appropriate File System *i*DataAgent installed.
  - On Windows there is no need to mount the snapshots that comprise the Recovery Point in order to back them up.
  - On UNIX, the snapshots do need to be mounted, but CDR does this automatically.
- The following types of backups are supported:
  - Full
  - Differential
  - Incremental

  See the following for step-by-step procedures:
  - Configure CDR for Backups of Recovery Points
  - Back up a Recovery Point
- Consider the following if you have configured Recovery Points to be backed up when they are created:
  - The creation of a new Recovery Point is not dependant on the successful completion of the backup of the previous Recovery Point. Therefore, if for some reason a backup job in this schedule is delayed/failed/pending/suspended, etc., Recovery Point creation will continue unaffected. When the **Maximum Number of Recovery Points** specified in the Replication Set Properties (Replication Options) tab is reached, CDR will begin deleting the oldest ones and it is thus possible that a Recovery Point will be pruned before the backup delay can be addressed. See Also: Maximum Number of Recovery Points.
  - When a backup job in this schedule is delayed/failed/pending/suspended, etc., it may be confusing as to which data has actually been backed up when the delayed backup job finally runs. It will back up the Recovery Point data as it existed when the backup initially started, even though the time of backup reflected by the backup history is later. Consider the following example:

    6:00a.m. - Recovery Point 1 is created, but backup fails.

    7:00a.m. - Recovery Point 2 is created, and backup completes successfully.

    8:00a.m. - Recovery Point 3 is created, and backup completes successfully.

    8:30a.m. - Administrator restarts the backup job that is in pending at 6:00a.m; job completes successfully, but job history shows a time of 8:30a.m., even though this is a backup of the point-in-time data from 6:00a.m.
- When backing up a Recovery Point on demand, ensure that you are aware of the age of the data that you are backing up. Each Recovery Point represents a specific point-in-time on the Destination computer, and each Consistent Recovery Point represents a specific point-in-time on the Source computer, but you can choose to back up any of them at any time. The time stamp on the backup will not match the point-in-time represented by the Recovery Point. As an example, let's look at a case where you have 12 Recovery Points, each created an hour apart starting with RP1 at 1:00 a.m., RP2 created at 2:00 a.m., and so on, and you create the following backups on-demand:
  - At 10:00 a.m. you back up RP9 (which was created at 9:00 a.m.) and label it "backup10am".
  - At 11:00 a.m. you back up RP2 (which was created at 2:00 a.m.) and label it "backup11am".

  Several days later you browse these two backups, and note the time of each backup as 10:00 a.m. and 11:00 a.m., but might not realize that the data in "backup11am" is actually older than the data in "backup10am", based upon when each Recovery Point was created. Thus, for on-demand backups of Recovery Points, you should be careful that you know the relative age of the data in each backup, not merely the time stamp of the backup itself, so that the correct data is selected for a restore operation.

RESTORE RECOVERY POINTS

The backup of Recovery Points or Consistent Recovery Points is performed using the Windows File System *i*DataAgent, which is also used to perform restore operations. For more information about restoring from a backup, see Restore Data - Windows File Systems.

DELETE RECOVERY POINTS

When the maximum number of Recovery Points is reached, they are automatically deleted by the system in the order in which it was created, starting with the oldest first. If one of the Recovery Points to be deleted is mounted, shared or in use, the following actions will be performed.

- On Windows, the system will delete next available for deletion Recovery Point. If there are no available Recovery Points for deletion the system will stop deleting Recovery Points until the situation is resolved.
- On UNIX, the system will mark the Recovery Point to be deleted when it is no longer in use.

The maximum number is either the maximum number you specified in the **Replication Set Properties**, or you can configure the maximum number of

Recovery Points that exist at any one time (per Replication Set for Windows, or per Replication Pair for UNIX) up to the following system limits:

- AIX with LVM: 15
- Linux with LVM: 32
- Linux with QSnap: 32
- Solaris with QSnap: 32
- Windows with VSS: 32 (256 for Fan-In configurations)
- Windows with ONTAP: 32 (256 for Fan-In configurations)

Obviously, the higher the limit you set, the greater will be the need for hard disk space.

Recovery Points can also be deleted manually, as long as they are unmounted, and for Windows, unshared. The system behavior is slightly different depending on the operating system:

- For Windows, Recovery Points created by VSS or ONTAP can be deleted in any order.
- For AIX, Recovery Points can be deleted in any order, but when deleted out-of-order, even though they will no longer appear when you browse in the CommCell Console, the snapshots that comprise a deleted Recovery Point that was not the oldest one will actually continue to exist until all older Recovery Points have been deleted.
- For Linux, Recovery Points can be deleted in any order, and the snapshots are deleted immediately.

For step-by-step procedures, see Delete a Recovery Point.

RECOVERING DATA FROM RECOVERY POINTS AND BACKUPS OF RECOVERY POINTS

Recovering data from Recovery Points or backups of Recovery Points is similar to Recovery Replicated Data. No additional or special steps are required. For more information on recovering data, see Recover Replicated Data.

IMPORTANT CONSIDERATIONS

Consider the following when using either type of Recovery Point with CDR.

GENERAL

- For Windows, Recovery Points will only be created for a Replication Set when all its Replication Pairs are in the "Replicating" Job State. If any Replication Pair is not in the "Replicating" state, no Recovery Point snapshots will be created for that Replication Set.

  For Fan-In, Recovery Points can be created for replication pair being in any Job State.

  For UNIX, any Replication Pair not in the "Replicating" state will be ignored and Recovery Point snapshots will be created for all Replication Pairs that are in the "Replicating" state.
- Persistence, the ability for the snapshot's integrity to be maintained, no matter the type of shutdown or reboot, is always enabled for Recovery Points.
- If cluster is used for Recovery Points creation, cluster resources should have persistent binding on.

CONSISTENT RECOVERY POINT FOR EXCHANGE DATA

- A Consistent Recovery Point can be successfully created for Exchange data, even if the Exchanges Stores are dismounted. If the Stores have been dismounted by the Exchange Server because of corruption, the Consistent Recovery Point will be created anyway, and will contain the corrupt data. Thus, it is essential to be aware of the state of any Stores that have been dismounted at the time a Consistent Recovery Point is created.

REPLICATION PAIRS FOR ORACLE

If Replications Pairs for Oracle log and database replication are configured to use the same destination volume, the actual number of Recovery Points retained will be the number specified in **Maximum Number of Recovery Points** divided by two, because separate snapshots will be created for the logs and the database. For example, if you have specified 10 as the maximum number of Recovery Points, but use the same destination volume for the logs and database, only 5 Recovery Points will be retained, with 2 snapshots in each.

SNAPSHOT SPACE REQUIREMENTS

It is essential to ensure that the destination volume specified for each Replication Pair has sufficient space for the maximum number of snapshots that will be created and retained.

For VSS, the cache can be configured using the **`vssadmin add shadowstorage`** command from a command line prompt. Refer to Microsoft documentation for details. Note that if the specified volume runs out of space, VSS will delete existing snapshots to make room for the new snapshots, even if your specified maximum number of Recovery Points has not been reached.

Back to Top

Recovery Points - How To

Topics | How To | Related Topics

CONFIGURE RECOVERY POINTS

Required Capability: Capabilities and Permitted Actions

Before you begin

- For Consistent Recovery Points involving supported application, you must Configure CDR to Replicate Application Data.

To configure Recovery Points:

1. From the CommCell Browser, right-click a Replication Set and select **Properties**.
2. In the Replication Set Properties (Replication Options) tab, select **Allow Recovery Points**.
3. In the **Maximum Number of Recovery Points** field, specify the maximum number to retain.

   
   Ensure that the destination volume specified for each Replication Pair has sufficient space for the maximum number of snapshots that will be created and retained.
4. On Windows you can configure to automatically mount the snapshots from the latest Recovery Point when they are created. (This is not available for ONTAP Snapshots.) On UNIX, you can configure automatically mounts every snapshot by default, and automatically create a symbolic link to the mount points it creates.
5. Perform the following steps to mount the snapshot:

   Select the **Auto-Mount snapshots after Recovery Point Creation** option, then click **Enter Mount Points**.

   In the Enter Mount Points dialog box, click the **Mount Point** column for each Destination Path and do one of the following:

   On Windows, do one of the following for each snapshot to be mounted:
   - Specify a drive letter (e.g., G:) to which to mount the snapshot. Ensure that you specify an unused drive letter.
   - Type a mount point name (e.g., G:\mountpoint1) to which the snapshot will be mounted.

   On UNIX, type a name for the symbolic link to the mount point that will be created by the system.

   Even though you specify the mount point for a recovery point on a UNC path, this mount point is ignored when a

recovery point is created on a UNC path.

6. For Windows, select a Snap Engine Type.

   For UNIX, CDR software automatically recognizes the File Systems configured on the destination computer and detects the appropriate snap engine. For more information on supported snap engines, see Snapshot Engines - Support.

7. In the **Pre/Post Process** tab, specify any Pre/Post Recovery Point commands to be run. Refer to Configure Pre/Post Processing.
8. If you will be backing up Recovery Points, see Configure CDR for Backups of Recovery Points.
9. Click **OK** to save your changes.

It is recommended that you also see Space Check for the Quick Recovery and ContinuousDataReplicator Agents and configure the Disk Space Low alert to provide warning that the destination computer, and the cache volume in particular, is running out of disk space, which will ultimately cause all recovery points to be deleted.

You can now Create a Recovery Point.

CREATE A RECOVERY POINT

Before You Begin

- Recovery Points cannot be created until you Configure Recovery Points.

Required Capability: See Capabilities and Permitted Actions

To create a Recovery Point:

1. From the CommCell Browser, right-click a Replication Set, and select **Create Recovery Points**.
2. From the Create Recovery Point dialog box, select the type, and whether to create the Recovery Point immediately or schedule it.
3. In addition, you can select the option of backing up Recovery Points. If you select this option, click **Enter Backup Details** to select the type of backup as well as any **Advanced Options**.
4. When you are finished, click **OK**.
5. If you opted to schedule Recovery Point creation, the **Schedule Details** dialog box will appear. Specify the details for your schedule; click **Options** for additional choices. When you are satisfied, click **OK**.

CREATE A RECOVERY POINT IN A FAN-IN CONFIGURATION

Before You Begin

- Review Using ContinuousDataReplicator in a Fan-In Configuration
- Recovery Points cannot be created until you Configure Recovery Points. Note that for Workstation Backup Agent, Recovery Point creation is enabled by default.

Required Capability: See Capabilities and Permitted Actions

To create a Recovery Point for a Fan-In configuration:

1. From the CommCell Browser, right-click **ContinuousDataReplicator** on the destination client, and select **Browse Destination Volumes**.
2. In the **Browse Destination Volumes** dialog box, right-click a destination volume, and select **Create Recovery Point**.
3. From the Create Recovery Point dialog box, select whether to create the Recovery Point immediately or schedule it.
4. In addition, you can select the option of backing up Recovery Points. If you select this option, click **Enter Backup Details** to select the type of backup as well as any Advanced Options.
5. When you are finished, click **OK**.
6. If you opted to schedule Recovery Point creation, the **Schedule Details** dialog box will appear. Specify the details for your schedule, click **Options** for additional choices. When you are satisfied, click **OK**.

DELETE A RECOVERY POINT

Before You Begin

- You cannot delete a Recovery Point that is currently mounted or shared. See Unmount or Unshare a Recovery Point.
- When Recovery Points are deleted manually, there are some restrictions and system behavior differences depending on the operating system; for more information, see Deleting Recovery Points.
- Deleting a Recovery Point will remove its corresponding snapshots as well as the Recovery Point Creation History.

Required Capability: See Capabilities and Permitted Actions

To delete a Recovery Point:

1. From the CommCell Browser, right-click a Replication Set, and select **Browse Recovery Points**.
2. From the **Recovery Points** dialog box, select a Recovery Point and click **Delete**.

DELETE A FAN-IN RECOVERY POINT

Before You Begin

- Review Using ContinuousDataReplicator in a Fan-In Configuration
- Deleting a Recovery Point will remove its corresponding snapshots as well as the Recovery Point Creation History.

Required Capability: See Capabilities and Permitted Actions

To delete a Fan-In Recovery Point:

1. From the CommCell Browser, right-click **ContinuousDataReplicator** on the destination client, and select **Browse Destination Volumes**.
2. In the **Browse Destination Volumes** dialog box, right-click a Recovery Point, and select **Delete**.

AUTOMATICALLY MOUNT A RECOVERY POINT

Before You Begin

- Review Mounting or Sharing Recovery Point Snapshots.

Required Capability: See Capabilities and Permitted Actions

To configure a Replication Set so that its Recovery Points are automatically mounted:

1. From the CommCell Browser, right-click a **Replication Set** for which at least one Replication Pair has already been created, and select **Properties**.
2. From the **Replication Options** tab of the Replication Set Properties screen, select the **Auto-Mount snapshots after Recovery Point Creation** option.
3. Click **Enter Mountpoints**.
4. In the Enter Mount Points dialog box, click the **Mount Point** column for each Destination Path and do one of the following:
   - For Windows, do one of the following for each snapshot to be mounted:

     Specify a drive letter (e.g., `G:`) to which the snapshot will be mounted. Ensure that you specify an *unused* drive letter.

     Type a mount point name (e.g., `G:\mountpoint1`) to which the snapshot will be mounted.
   - For UNIX, type a name for the symbolic link to the mount point created by the system.
5. Click **OK**.

MANUALLY MOUNT OR SHARE THE SNAPSHOTS OF AN EXISTING RECOVERY POINT

Before You Begin

- Review Mounting or Sharing Recovery Point Snapshots.
- For UNIX, mount points are automatically created for all Recovery Point snapshots, and they are automatically mounted.

Required Capability: See Capabilities and Permitted Actions

To mount or share the snapshots of an existing Recovery Point, on Windows:

1. In the CommCell Browser, perform a Browse operation for one of the following:
   - Replication Set
   - CDR Agent
   - Client
   - Application *i*DataAgent (Exchange, SQL, Oracle)

2. Do one of the following, depending on the object you browsed:
   - From the **Recovery Points** dialog box, select a Recovery Point and click **View Snapshots**.

     In the Snapshots dialog box, select a snapshot and click **Mount**.
   - In the Browse dialog box, right-click the snapshot to be mounted, and select **Mount Snapshot**.
3. Do one of the following:
   - specify a drive letter (e.g., `G:`) to which to mount the Recovery Point. Ensure that you specify an *unused* drive letter.
   - type a mount point name (e.g., `X:\mountpoint1`) for the Recovery Point.

   
   For Windows 2003, Recovery Point snapshots created with VSS, cannot be mounted on a mount point (e.g. `X:\mountpoint1`), this needs to be mounted to a drive letter (e.g., `G:`).
4. If you want to create a share, do one of the following, depending on which object you browsed:
   - In the Snapshots dialog box, select a mounted snapshot and click **Share**.
   - In the **Browse** dialog box, right-click the mounted snapshot to be shared, and select **Create Network Share**.

   In the Enter Share Name dialog box, type a name for the share and select **Make share writeable** option. Click **OK**.
5. Close the **Snapshots** dialog box or **Browse** window.

AUTOMATICALLY MOUNT SNAPSHOTS FOR FAN-IN RECOVERY POINT

Before you Begin

- Review How to Setup Fan-In Configuration
- You can configure CDR so that the snapshots that comprise the latest Fan-In Recovery Point will be automatically mounted when it is created. This will affect all Replication Sets that use the computer as a Fan-In destination. Alternately, once Recovery Points have been created, you can select an existing Recovery Point and mount its snapshots.

Required Capability: See Capabilities and Permitted Actions

To configure CDR to automatically mount snapshots for Fan-In Recovery Points:

1. From the CommCell Browser, right-click **ContinuousDataReplicator** on the destination client, and select **Properties**.
2. From the **Fan-In** tab, select the **Recovery Points** tab, and select the **Auto-Mount snapshots after Recovery Point Creation** option.
3. Click **Enter Mountpoints**.
4. In the Enter Mount Points dialog box, click the **Mount Point** column for each Destination Path and specify a drive letter (e.g., `G:`) to which the snapshot will be mounted.

   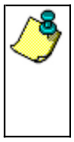
   - Ensure that you specify an *unused* drive letter.
   - Do not specify mount points for a UNC path. If you specify the mount point for a recovery point on a UNC path, this mount point is ignored when a recovery point is created on a UNC path.
   - When a Recovery Point snapshots created with ONTAP snap engine and **Auto-Mount snapshots after Recovery Point Creation** option is selected, share of snapshot will be created.
5. Click **OK**.

MANUALLY MOUNT OR SHARE THE SNAPSHOTS OF AN EXISTING FAN-IN RECOVERY POINT

Before you Begin

- Review How to Setup Fan-In Configuration

Required Capability: See Capabilities and Permitted Actions

To mount the snapshots of an existing Fan-In Recovery Point, on Windows:

1. From the CommCell Browser, right-click **ContinuousDataReplicator** on the destination client, and select **Browse Destination Volumes**.
2. In the **Browse Destination Volumes** dialog box, right-click a Recovery Point, and select **Mount**.
3. Specify a drive letter (e.g., `G:`) to which to mount the Recovery Point. Ensure that you specify an *unused* drive letter.
4. To share a snapshot on ONTAP, do the following, depending on which object you browsed:
   - In the **Browse** window, right-click the mounted snapshot to be shared, and select **Create Network Share**.

- In the Enter Share Name dialog box, type a name for the share.
- Select **Make share writeable** option. Click **OK**.

5. Close the **Browse** window.

UNSHARE THE SNAPSHOTS OF A RECOVERY POINT

Before You Begin

- Review Mounting or Sharing Recovery Point Snapshots.

Required Capability: See Capabilities and Permitted Actions

To unshare the snapshots of a Recovery Point (On Windows only):

1. In the CommCell Browser, perform a Browse operation for one of the following:
   - Replication Set
   - CDR Agent
   - Client
   - Application *i*DataAgent (Exchange, SQL, Oracle)
2. Do one of the following, depending on which object you browsed:
   - From the **Recovery Points** dialog box, select a Recovery Point and click **View Snapshots**.

     In the Snapshots dialog box, select a shared snapshot and click **Unshare**.
   - In the **Browse** dialog box, right-click a mounted snapshot, and select **Delete Network Share**.
3. Close the Snapshots dialog box or Browse dialog box.

UNMOUNT THE SNAPSHOTS OF A RECOVERY POINT

Before You Begin

- Review Mounting or Sharing Recovery Point Snapshots.

Required Capability: See Capabilities and Permitted Actions

To unmount the snapshots of a Recovery Point:

1. In the CommCell Browser, perform a Browse operation for one of the following:
   - Replication Set
   - CDR Agent
   - Client
   - Application *i*DataAgent (Exchange, SQL, Oracle)
2. Do one of the following, depending on which object you browsed:
   - From the **Recovery Points** dialog box, select a Recovery Point and click **View Snapshots**.

     In the Snapshots dialog box, select a mounted snapshot and click **Unmount**.
   - In the Browse dialog box, right-click a mounted snapshot, and select **Unmount Snapshot**.
3. Close the Snapshots dialog box or Browse dialog box.

UNMOUNT A FAN-IN RECOVERY POINT

Before You Begin

- Review Using ContinuousDataReplicator in a Fan-In Configuration.

Required Capability: See Capabilities and Permitted Actions

To unmount the snapshots of a Fan-In Recovery Point:

1. From the CommCell Browser, right-click **ContinuousDataReplicator** on the destination client, and select **Browse Destination Volumes**.
2. In the **Browse Destination Volumes** dialog box, right-click a Recovery Point, and select **Unmount**.
3. Close the **Browse Destination Volumes** dialog box.

CONFIGURE CDR FOR BACKUPS OF RECOVERY POINTS

Before You Begin

- Review Recovery Points - Backups of Recovery Points.

Required Capability: Capabilities and Permitted Actions

To configure CDR for backups of Recovery Points:

1. From the CommCell Browser, right-click a Replication Set, and select **Properties**. In the Replication Set Properties (Storage Policy) tab, select **Back up Recovery Points**, and specify a Storage Policy and Backup Set for the destination machine.
2. Click **OK** to save your changes.

- To perform backups of a Recovery Point, see Back up a Recovery.
- If you change the Storage Policy for a Replication Set, the next backup will be converted to a full.

CONFIGURE CDR FOR BACKUPS OF RECOVERY POINTS IN A FAN-IN CONFIGURATION

Before You Begin

- Review Using ContinuousDataReplicator in a Fan-In Configuration.

Required Capability: Capabilities and Permitted Actions

To configure CDR for backups of Recovery Points in a Fan-In configuration:

1. From the CommCell Browser, right-click **ContinuousDataReplicator** on the destination client, and select **Properties**.
2. From the Fan-In tab, select **Back up Recovery Points**, and specify a Storage Policy for the destination machine.
3. Click **OK** to save your changes.

- To perform backups of a Recovery Point, see Back up a Recovery Point.
- If you change the Storage Policy, the next backup will be converted to a full.

BACK UP A RECOVERY POINT

Before You Begin

- Review Recovery Points - Backups of Recovery Points.
- Before you can perform backups of Recovery Points, you must first Configure CDR for Backups of Recovery Points.
- You can either configure a Replication Set so that its Recovery Points are backed up when they are created, or you can select an existing Recovery Point and back it up on-demand.

Required Capability: See Capabilities and Permitted Actions

To configure Recovery Points to be backed up when they are created:

1. From the CommCell Browser, right-click a Replication Set, and select **Create Recovery Points**.
2. From the Create Recovery Point dialog box, make your selections for **Recovery Point Type** and **Job Initiation**, then select **Back up Recovery Point**.
3. Click **Enter Backup Details** to select the type of backup from the **Enter Backup Details** dialog box.
4. From the **Enter Backup Details** dialog box, click **Advanced** to display the **Advanced Backup Options** dialog box, from which you can select additional options for the backup operation.
5. When you are finished, click **OK** to exit each of the dialog boxes.

To perform a full backup of an existing Recovery Point on-demand:

1. From the CommCell Browser, right-click a Replication Set, and select **Browse Recovery Points**.
2. From the **Recovery Points** dialog box, select a Recovery Point and click **Backup**.
3. An immediate full backup is performed for each of the snapshots that comprise the selected Recovery Point; in addition, the next backup performed will be converted to a full backup.

CONFIGURE A REPLICATION SET FOR PRE/POST PROCESSES FOR RECOVERY POINT CREATION

Before You Begin

- Review Recovery Points.
- It is recommended to not configure a pre/post process for a Replication Set that is currently replicating data.
- Verify that there are no pre/post processes already assigned for the Replication Set.
- Pre-process commands 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 Replication Set for Pre/Post processes for Recovery Point creation:

1. In the CommCell Browser, right-click the Replication Set for which you want to configure a pre/post process, then click **Properties** from the shortcut menu.
2. In the Properties dialog box, click the Pre/Post Process tab.
3. 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).
   - Pre-Recovery Point Command
   - Post-Recovery Point Command
4. In the User Account dialog box, select **Use Local System Account,** or select **Impersonate User** and enter a user name and password. Click **OK**.
5. Click **OK** to save your changes and close the Pre/Post Process tab of the Properties dialog box.

SPECIFY A SNAP ENGINE

Before You Begin

- Review Snapshot Engines - Support
- In the source computer:
  - For Windows 2003 or higher, VSS is used as the default snap engine.
- This procedure only applies to the destination computer.

Required Capability: Capabilities and Permitted Actions

To specify a snap engine on Windows to the destination computer:

1. In the CommCell Browser, right-click the Replication Set on the source computer and select **Properties**.
2. In the Replication Options tab, select one of the choices in the **Select Snap Engine Type for Recovery Point Creation** section.

   For Windows, **VSS** is the default selection; **ONTAP** can be used if available on the destination computer where the Recovery Point will be created.

   For UNIX, Snap Engine Type is automatically detected by the File System configured on the destination computer.
3. Click **OK** to save your changes.

   On UNIX, CDR software automatically recognizes the File Systems configured on the destination computer and detects the appropriate snap engine. For more information on supported snap engines, see Snapshot Engines - Support.

CHANGE ACCOUNT FOR EXECUTING PRE/POST COMMANDS (RECOVERY POINTS)

Required Capability: See Capabilities and Permitted Actions

To change a user account for executing pre/post commands for creating Recovery Points on Windows:

1. From the CommCell Browser, expand the tree to view the Replication Set for the affected ContinuousDataReplicator agent.
2. Right-click the appropriate Replication Set and select **Properties**.
3. In the **Replication Set Properties** dialog box, select the **Pre/Post Process** tab.
4. From the **Pre/Post Process** tab, click **Change**.
5. From the User Account dialog box, select one of the account options. If you select **Impersonate User**, type the appropriate user name and password.
6. Click **OK** to save the settings.

Recover Replicated Data

Topics | How To | Related Topics

OVERVIEW

The data replicated by ContinuousDataReplicator (CDR) is copied from a source computer to a destination computer. In addition, Recovery Points can be created from snapshots of the data on the destination computer. Finally, Recovery Points can be backed up. Thus there are multiple ways in which your data can be protected using CDR, and several options for recovery.

- Data that has been replicated to the Destination computer, and is still residing there, can be recovered using Copyback.
- Recovery Points can also be recovered using Copyback.
- Recovery Points can be mounted on the destination computer or, on Windows, made available on the network as shared volumes, allowing the data to be directly copied.
- Data backed up from a Recovery Point can be restored in a normal manner using the Windows File System *i*DataAgent.

RECOVERY SCENARIOS

COPY BACK FROM RECOVERY POINT OR DESTINATION COMPUTER

From a Recovery Point on a Destination computer, or the Live Copy on the Destination computer itself, a Copyback operation can copy that data to a designated volume on any client.

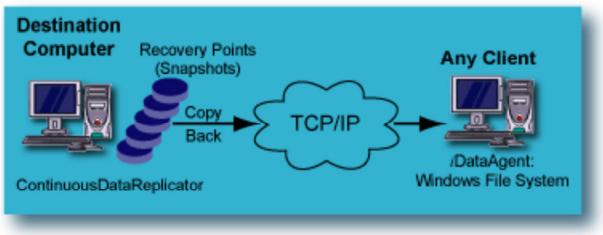

MOUNT OR SHARE A RECOVERY POINT ON ANY COMPUTER

Recovery Points created on the Destination computer can be mounted, or for Windows they can also be shared, and thus made available to any computer on the network, so that files can be copied from them.

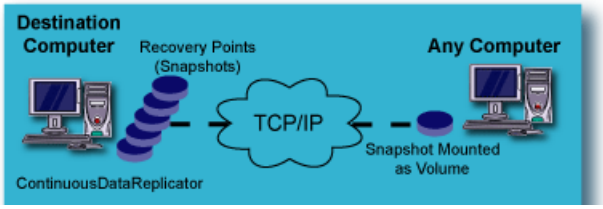

RESTORE DATA FROM BACKUP

File system data which has been backed up from a Recovery Point on a Destination computer, can be restored using a File System *i*DataAgent.

SQL and Exchange data which has been backed up from a Consistent Recovery Point on a Destination computer using the File System *i*DataAgent, can be browsed and restored using the Windows File System *i*DataAgent on the Client.

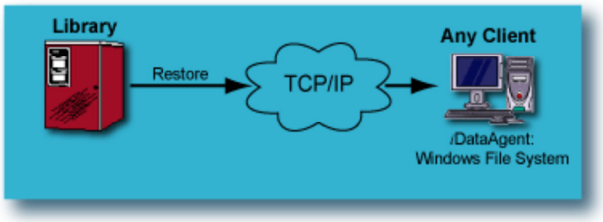

COPYBACK

A Copyback operation copies the contents of a destination path (Live Copy) or Recovery Point to a specified recovery host and recovery path. This is true for both file system data, as well as application data. The Live Copy or multiple Recovery Points can be copied back by selecting them during the browse, and then specifying the destination for the copyback operation. Copyback can be performed from any of the levels described in Browse Data - ContinuousDataReplicator.

For step-by-step instructions for file system data, see Copy Back File System Data from a Recovery Point or the Live Copy

COPYBACK OF EXCHANGE DATA

Exchange data is restored at the Storage Group level. While you can restore Exchange data from a Recovery Point, a backup of a Recovery Point, or the Live Copy, these methods will not ensure consistency of the application data; only a restore from a Consistent Recovery Point, or a backup of one, will ensure consistency of application data. Copyback is recommended as the primary method of moving the replicated data back to the production Exchange Server, in addition to restoring a backup of a Consistent Recovery Point where that is appropriate. Exchange circular logging can be either enabled or disabled, with no impact on the Copyback operation. To ensure application integrity, you must use **Add App** to create your Replication Pairs. (Refer to Add or Edit a Replication Pair and Application Integration.) **Add App** discovers the location of the Exchange `.chk` file and `tmp.edb` file, which means the Exchange `.chk` file will not have to be deleted before performing the Copyback operation, since it is restores from the same point-in-time as the database and log files.

For step-by-step instructions, see Copy Back Exchange Data from a Consistent Recovery Point.

COPYBACK OF SQL DATA

SQL data is restored at the database level. While you can restore SQL data from a Recovery Point, a backup of a Recovery Point, or the Live Copy, these methods will not ensure consistency of the application data; only a restore from a Consistent Recovery Point, or a backup of one, will ensure consistency of application data. Copyback is recommended as the primary method of moving the replicated data back to the production SQL Server, in addition to restoring a backup of a Consistent Recovery Point where that is appropriate. To ensure application integrity, you must use **Add App** to create your Replication Pairs. (Refer to Add or Edit a Replication Pair and Application Integration.) **Add App** discovers the location of, not only user-defined databases (`.mdf`, `.ndf`, `.ldf`) but also any system databases on the client, which you can select for data replication.

For step-by-step instructions, see Copy Back SQL Data from a Consistent Recovery Point.

COPYBACK OF ORACLE DATA

Oracle data is restored at the database level. While you can restore Oracle data from a Recovery Point, a backup of a Recovery Point, or the Live Copy, these methods will not ensure consistency of the application data; only a restore from a Consistent Recovery Point, or a backup of one, will ensure consistency of application data. A Consistent Recovery Point is recommended as the source when moving replicated data back to the production Oracle server. To ensure application integrity, you must use **Add App** to create your Replication Pairs. (Refer to Add or Edit a Replication Pair and Application Integration.) **Add App** discovers the location of all user-defined and system databases on the client, which you can select for data replication.

For step-by-step instructions, see Copy Back Oracle Data from a Consistent Recovery Point.

For Unix, sparse files attributes are not retained during Copyback, neither from a Live Copy nor a Recovery Point; the files assume the attributes of regular files on the recovery host.

RECOVERY DESTINATIONS

CDR can perform a Copyback operation either to the client computer from which it originated, comparable to an in-place restore on other agents, or to a different computer, comparable to an out-of-place restore.

You can also mount or share the Recovery Points, providing access to other computers on the network. Keep in mind that mounting or creating a share of a Recovery Point is not intended to recover data into a live production environment; applications will not be quiesced. Volumes will not be ready for use by their associated applications. This only provides an alternative means to access data in the Recovery Points.

CDR data can also be restored to any client computer in the same CommCell with Windows File System Agent installed.

In addition, data that has been backed up can be restored in the usual manner.

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

IN-PLACE RECOVERY

For both file system and application data, copyback from either the live copy or from Recovery Points, restore from a backup of a Recovery Point, and mounting or sharing of Recovery Points are all suitable methods of restoring data to the location from which it originated.

OUT-OF-PLACE RECOVERY

For file system data only, copyback from either the live copy or from Recovery Points, restore from a backup of a Recovery Point, and mounting or sharing of Recovery Points are all suitable methods of restoring data to a different location than the one from which it originated.

CROSS-PLATFORM RECOVERY

For file system data only, restore from a backup of a Recovery Point, and mounting or sharing of Recovery Points are all suitable methods of restoring data cross-platform. However, see Restore Destinations and Cross-Platform Restores for more information and considerations.

RECOVERY CONSIDERATIONS FOR THIS AGENT

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

- For an In-Place recovery of application data, the application should be down or off-line during the restore.

Back to Top

Recover Replicated Data - How To

Topics | How To | Related Topics

AUTOMATICALLY MOUNT A RECOVERY POINT

Before You Begin

- Review Mounting or Sharing Recovery Point Snapshots.

Required Capability: See Capabilities and Permitted Actions

To configure a Replication Set so that its Recovery Points are automatically mounted:

1. From the CommCell Browser, right-click a **Replication Set** for which at least one Replication Pair has already been created, and select **Properties**.
2. From the **Replication Options** tab of the Replication Set Properties screen, select the **Auto-Mount snapshots after Recovery Point Creation** option.
3. Click **Enter Mountpoints**.

4. In the Enter Mount Points dialog box, click the **Mount Point** column for each Destination Path and do one of the following:
   - For Windows, do one of the following for each snapshot to be mounted:

     Specify a drive letter (e.g., `G:`) to which the snapshot will be mounted. Ensure that you specify an *unused* drive letter.

     Type a mount point name (e.g., `G:\mountpoint1`) to which the snapshot will be mounted.
   - For UNIX, type a name for the symbolic link to the mount point created by the system.
5. Click **OK**.

MANUALLY MOUNT OR SHARE THE SNAPSHOTS OF AN EXISTING RECOVERY POINT

Before You Begin

- Review Mounting or Sharing Recovery Point Snapshots.
- For UNIX, mount points are automatically created for all Recovery Point snapshots, and they are automatically mounted.

Required Capability: See Capabilities and Permitted Actions

To mount or share the snapshots of an existing Recovery Point, on Windows:

1. In the CommCell Browser, perform a Browse operation for one of the following:
   - Replication Set
   - CDR Agent
   - Client
   - Application *i*DataAgent (Exchange, SQL, Oracle)
2. Do one of the following, depending on the object you browsed:
   - From the **Recovery Points** dialog box, select a Recovery Point and click **View Snapshots**.

     In the Snapshots dialog box, select a snapshot and click **Mount**.
   - In the Browse dialog box, right-click the snapshot to be mounted, and select **Mount Snapshot**.
3. Do one of the following:
   - specify a drive letter (e.g., `G:`) to which to mount the Recovery Point. Ensure that you specify an *unused* drive letter.
   - type a mount point name (e.g., `X:\mountpoint1`) for the Recovery Point.

   
   For Windows 2003, Recovery Point snapshots created with VSS, cannot be mounted on a mount point (e.g. `X:\mountpoint1`), this needs to be mounted to a drive letter (e.g., `G:`).
4. If you want to create a share, do one of the following, depending on which object you browsed:
   - In the Snapshots dialog box, select a mounted snapshot and click **Share**.
   - In the **Browse** dialog box, right-click the mounted snapshot to be shared, and select **Create Network Share**.

   In the Enter Share Name dialog box, type a name for the share and select **Make share writeable** option. Click **OK**.
5. Close the **Snapshots** dialog box or **Browse** window.

AUTOMATICALLY MOUNT SNAPSHOTS FOR FAN-IN RECOVERY POINT

Before you Begin

- Review How to Setup Fan-In Configuration
- You can configure CDR so that the snapshots that comprise the latest Fan-In Recovery Point will be automatically mounted when it is created. This will affect all Replication Sets that use the computer as a Fan-In destination. Alternately, once Recovery Points have been created, you can select an existing Recovery Point and mount its snapshots.

Required Capability: See Capabilities and Permitted Actions

To configure CDR to automatically mount snapshots for Fan-In Recovery Points:

1. From the CommCell Browser, right-click **ContinuousDataReplicator** on the destination client, and select **Properties**.
2. From the **Fan-In** tab, select the **Recovery Points** tab, and select the **Auto-Mount snapshots after Recovery Point Creation** option.
3. Click **Enter Mountpoints**.
4. In the Enter Mount Points dialog box, click the **Mount Point** column for each Destination Path and specify a drive letter (e.g., `G:`) to which the snapshot

will be mounted.

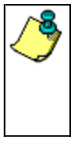

- Ensure that you specify an *unused* drive letter.
- Do not specify mount points for a UNC path. If you specify the mount point for a recovery point on a UNC path, this mount point is ignored when a recovery point is created on a UNC path.
- When a Recovery Point snapshots created with ONTAP snap engine and **Auto-Mount snapshots after Recovery Point Creation** option is selected, share of snapshot will be created.

5. Click **OK**.

MANUALLY MOUNT OR SHARE THE SNAPSHOTS OF AN EXISTING FAN-IN RECOVERY POINT

Before you Begin

- Review How to Setup Fan-In Configuration

Required Capability: See Capabilities and Permitted Actions

To mount the snapshots of an existing Fan-In Recovery Point, on Windows:

1. From the CommCell Browser, right-click **ContinuousDataReplicator** on the destination client, and select **Browse Destination Volumes**.
2. In the **Browse Destination Volumes** dialog box, right-click a Recovery Point, and select **Mount**.
3. Specify a drive letter (e.g., `G:`) to which to mount the Recovery Point. Ensure that you specify an *unused* drive letter.
4. To share a snapshot on ONTAP, do the following, depending on which object you browsed:
   - In the **Browse** window, right-click the mounted snapshot to be shared, and select **Create Network Share**.
   - In the Enter Share Name dialog box, type a name for the share.
   - Select **Make share writeable** option. Click **OK**.
5. Close the **Browse** window.

UNSHARE THE SNAPSHOTS OF A RECOVERY POINT

Before You Begin

- Review Mounting or Sharing Recovery Point Snapshots.

Required Capability: See Capabilities and Permitted Actions

To unshare the snapshots of a Recovery Point (On Windows only):

1. In the CommCell Browser, perform a Browse operation for one of the following:
   - Replication Set
   - CDR Agent
   - Client
   - Application *i*DataAgent (Exchange, SQL, Oracle)
2. Do one of the following, depending on which object you browsed:
   - From the **Recovery Points** dialog box, select a Recovery Point and click **View Snapshots**.

     In the Snapshots dialog box, select a shared snapshot and click **Unshare**.
   - In the **Browse** dialog box, right-click a mounted snapshot, and select **Delete Network Share**.
3. Close the Snapshots dialog box or Browse dialog box.

UNMOUNT THE SNAPSHOTS OF A RECOVERY POINT

Before You Begin

- Review Mounting or Sharing Recovery Point Snapshots.

Required Capability: See Capabilities and Permitted Actions

To unmount the snapshots of a Recovery Point:

1. In the CommCell Browser, perform a Browse operation for one of the following:
   - Replication Set

- CDR Agent
- Client
- Application iDataAgent (Exchange, SQL, Oracle)

2. Do one of the following, depending on which object you browsed:
   - From the **Recovery Points** dialog box, select a Recovery Point and click **View Snapshots**.

     In the Snapshots dialog box, select a mounted snapshot and click **Unmount**.
   - In the Browse dialog box, right-click a mounted snapshot, and select **Unmount Snapshot**.
3. Close the Snapshots dialog box or Browse dialog box.

UNMOUNT A FAN-IN RECOVERY POINT

Before You Begin

- Review Using ContinuousDataReplicator in a Fan-In Configuration.

Required Capability: See Capabilities and Permitted Actions

To unmount the snapshots of a Fan-In Recovery Point:

1. From the CommCell Browser, right-click **ContinuousDataReplicator** on the destination client, and select **Browse Destination Volumes**.
2. In the **Browse Destination Volumes** dialog box, right-click a Recovery Point, and select **Unmount**.
3. Close the **Browse Destination Volumes** dialog box.

COPYBACK FILE SYSTEM DATA FROM A RECOVERY POINT OR THE LIVE COPY

Before You Begin:

- Review Copyback

Required Capability: See Capabilities and Permitted Actions

To Copyback file system data from a Recovery Point or the Live Copy:

1. In the CommCell Browser, perform a Browse operation for one of the following:
   - Replication Set
   - CDR Agent
   - Client
2. Do one of the following, depending on which object you browsed:
   - In the Recovery Points dialog box, select the Live Copy and/or Recovery Point(s) to be copied back, and click **Copyback**.
   - In the Browse dialog box, select the Live Copy and/or Recovery Point(s) to be copied back, and click **Recover all Selected...**
3. In the Copyback dialog box, select a **Recovery Host**, the host to which the data will be copied back.
4. Specify a **Recovery Path** for the live data and/or Recovery Point(s) you want to copy back. The Recovery Path must be larger than the data being copied back. To assign a **Recovery Path**, do one of the following:
   - double-click the path listed in the **Recovery Path** column and type a path on the selected host
   - click the **Browse 'Recovery Host'** button and browse to a destination path on the selected host
5. Optionally, select **Overwrite existing data under 'Recovery Path'**.
6. Click **OK** to copy back the data. The system displays a progress bar as it copies back the data. You can track the progress of the recover operation from the Job Controller dialog box.

COPYBACK FILE SYSTEM DATA FROM A FAN-IN RECOVERY POINT

Before You Begin

- Review Copyback
- Review Using ContinuousDataReplicator in a Fan-In Configuration

Required Capability: See Capabilities and Permitted Actions

To Copyback file system data from a Fan-In Recovery Point:

1. From the CommCell Browser, right-click **ContinuousDataReplicator** on the destination client, and select **Browse Destination Volumes**.
2. In the **Browse Destination Volumes** dialog box, right-click a Recovery Point, and select **Copyback**.
3. In the Copyback dialog box, select a **Recovery Host**, the host to which the data will be copied back.
4. Specify a **Recovery Path** for the Recovery Point(s) you want to copy back. The Recovery Path must be larger than the data being copied back. To assign a **Recovery Path**, do one of the following:
   - double-click the path listed in the **Recovery Path** column and type a path on the selected host
   - click the **Browse 'Recovery Host'** button and browse to a destination path on the selected host
5. Optionally, select **Overwrite existing data under 'Recovery Path'**.
6. Click **OK** to copy back the data. The system displays a progress bar as it copies back the data. You can track the progress of the recover operation from the Job Controller dialog box.

COPYBACK EXCHANGE DATA FROM A CONSISTENT RECOVERY POINT

Before You Begin

- Review Copyback
- This procedure assumes that you used **Add App** to create your Replication Pairs. For more information, refer to Add a Replication Pair and Application Integration.
- While it is possible to copy back Exchange data from the Live Copy, this data is not in a consistent state and thus application integrity is not ensured.

Required Capability: See Capabilities and Permitted Actions

To Copyback Exchange data from a Consistent Recovery Point:

1. For Outlook 2003, change the profile for any mailboxes being restored which have "Cached Mode" enabled, to **DISABLE CACHED MODE**. "Cached mode" can be enabled again after the Copyback operation completes successfully.
2. Abort the Replication Pairs for the source volume to which you will copy back Exchange data. This avoids slow copyback operation on the source volume.
3. Dismount the Exchange Stores for the Storage Group to which you will copy back point-in-time data.
4. Delete or rename all Exchange transaction logs on the Production Exchange Server for the Storage Group being restored.
5. In the CommCell Browser, perform a Browse operation for one of the following:
   - Replication Set
   - CDR Agent
   - Client
   - Application *i*DataAgent (Exchange)
6. Do one of the following, depending on which object you browsed:
   - In the Recovery Points dialog box, select the Consistent Recovery Point(s) to be copied back, and click **Copyback**.
   - In the Browse dialog box, select the Consistent Recovery Point(s) to be copied back, and click **Recover all Selected...**
7. In the Copyback dialog box, select a **Recovery Host**, the host to which the data will be copied back.
8. Specify a **Recovery Path** for the data you want to copy back. The Recovery Path must be larger than the data being copied back. To assign a **Recovery Path**, do one of the following:
   - double-click the path listed in the **Recovery Path** column and type a path on the selected host
   - click the **Browse 'Recovery Host'** button and browse to a destination path on the selected host
9. Select **Overwrite existing data under 'Recovery Path'**.
10. Click **OK** to copy back the data. The system displays a progress bar as it copies back the data. You can track the progress of the recover operation from the **Job Controller** dialog box.
11. Mount the Exchange Stores.
12. For Outlook 2003, the profile for any mailbox(es) changed to **DISABLE CACHED MODE** can be changed back to Cached Mode. The user will be able to see the recovered data instead of the cached copy.

COPYBACK SQL DATA FROM A CONSISTENT RECOVERY POINT

Before You Begin

- Review Copyback
- This procedure assumes that you used **Add App** to create your Replication Pairs. For more information, refer to Add or Edit a Replication Pair and Application Integration.
- While it is possible to copy back SQL data from the Live Copy, this data is not in a consistent state and thus application integrity is not ensured.

Required Capability: See Capabilities and Permitted Actions

To Copyback SQL data from a Consistent Recovery Point:

1. Abort the Replication Pairs for the source volume to which you will copy back SQL data. This avoids slow copyback operation on the source volume.
2. Detach the SQL database(s) for which you will copy back point-in-time SQL data.
3. In the CommCell Browser, perform a Browse operation for one of the following:
   - Replication Set
   - CDR Agent
   - Client
   - Application *i*DataAgent (SQL)
4. Do one of the following, depending on which object you browsed:
   - In the Recovery Points dialog box, select the Consistent Recovery Point(s) to be copied back, and click **Copyback**.
   - In the Browse dialog box, select the Consistent Recovery Point(s) to be copied back, and click **Recover all Selected...**
5. In the Copyback dialog box, select a **Recovery Host**, the host to which the data will be copied back.
6. Specify a **Recovery Path** for the data you want to copy back. The Recovery Path must be larger than the data being copied back. To assign a **Recovery Path**, do one of the following:
   - double-click the path listed in the **Recovery Path** column and type a path on the selected host
   - click the **Browse 'Recovery Host'** button and browse to a destination path on the selected host
7. Select **Overwrite existing data under 'Recovery Path'**.
8. Click **OK** to copy back the data. The system displays a progress bar as it copies back the data. You can track the progress of the recover operation from the Job Controller dialog box.
9. Attach your SQL database(s).

COPYBACK ORACLE DATA FROM A CONSISTENT RECOVERY POINT

Before You Begin

- Review Copyback
- Review Using ContinuousDataReplicator with Oracle
- This procedure assumes that you used **Add App** to create your Replication Pairs. For more information, refer to Add or Edit a Replication Pair and Application Integration.
- While it is possible to copy back Oracle data from the Live Copy, this data is not in a consistent state and thus application integrity is not ensured.
- Each Consistent Recovery Point for a given Oracle database comprises two snapshots, which will be referred to in the procedures as "first snapshot" and "second snapshot" according to the following definition of content:
  - The first snapshot includes all Oracle data volumes.
  - The second snapshot includes all Oracle log volumes and backup control file for the database at the time the Consistent Recovery Point was created.

Required Capability: See Capabilities and Permitted Actions

To Copyback Oracle data from a Consistent Recovery Point:

1. Abort the Replication Pairs for the source volume to which you will copy back Oracle data. This avoids slow copyback operation on the source volume.
2. Shut down the Oracle database.
3. Copy back the following Oracle data from the second snapshot (the snapshot of any destination log volume) in the Consistent Recovery Point you have selected, to a temporary directory:
   - `Backup.ctl.galaxy`
   - `dbconf.cfg`
   - `GalaxyControlFile.conf`
   - `init<instancename>.ora`

- SPFILE<*instancename*>.ora

4. This can be done using one of the following methods:
   - Mount the Recovery Point, then copy only the files listed above to a temporary directory.
   - Perform a Copyback operation. Since this will *not* allow you to select which files to copy back, only use this method if you want to copy back everything (database, logs, and files listed above) - assuming the original folders are empty. Proceed as follows:
   - In the CommCell Browser, perform a Browse operation for one of the following:

     Replication Set

     DR Agent

     Client

     Application *i*DataAgent (Oracle)
   - Do one of the following, depending on which object you browsed:

     In the Recovery Points dialog box, select the Consistent Recovery Point(s) to be copied back, and click **Copyback**.

     In the Browse dialog box, select the Consistent Recovery Point(s) to be copied back, and click **Recover all Selected...**
   - In the Copyback dialog box, select a **Recovery Host**, the host to which the data will be copied back.
   - Specify a **Recovery Path** for the data you want to copy back. The Recovery Path must be larger than the data being copied back. To assign a **Recovery Path**, do one of the following:

     Double-click the path listed in the **Recovery Path** column and type a path on the selected host

     Click the **Browse 'Recovery Host'** button and browse to a destination path on the selected host
   - Select **Overwrite existing data under 'Recovery Path'**.
   - Click **OK** to copy back the data. The system displays a progress bar as it copies back the data. You can track the progress of the recover operation from the Job Controller dialog box.
5. Create the database tracing directories (**udump**, **bdump**, **cdump**, etc.) for the Oracle database according to the dbconf.cfg file you copied back in Step 3. (<*Snap2Dir*>/archivelog_dest1/dbconf.cfg)
6. If the source database uses spfile, copy spfile from the location you copied back to in Step 3 to the default location under ORACLE_HOME. For example:

   cp <*Snap2Dir*>/archivelog_dest1/spfile <ORACLE SID>.ora <*default spfile file path*>
7. If the source database uses pfile, copy pfile from the location you copied back to in Step 3 to the default location under ORACLE_HOME. For example:

   cp <*Snap2Dir*>/archivelog_dest1/pfile <ORACLE SID>.ora <*default pfile file path*>
8. Copy the backup control file from the first archive log destination (from the second snapshot) to all three control files destinations. For example:

   cp <*Snap2Dir*>/archivelog_dest1/backup.ctl.galaxy <*controlfile1_Path from spfile*>/ control01.ctl
   cp <*Snap2Dir*>/archivelog_dest1/backup.ctl.galaxy <*Controlfile2_Path from spfile*>/ control02.ctl
   cp <*Snap2Dir*>/archivelog_dest1/backup.ctl.galaxy <*Controlfile3_Path from spfile*>/ control03.ctl

   
   GalaxyControlFile.conf also contains the locations of the control files at the time the Consistent Recovery Point was created.
9. If you did not perform a Copyback operation in Step 3, copy all data and log files to the exact same path as that of the original database, either by mounting the Recovery Point and copying the files, or by performing a Copyback operation. In case the same destination volume is used for database and logs, mount the first snapshot and copy all Oracle data files from the first snapshot to the path that the control file expects. For example:

   cp -a <*Snap1Dir*>/oradata/rman1/*.dbf /oradata/rman1/
10. Copy all the log files for all destinations from the second snapshot to the paths that the control file expects, either by mounting the Recovery Point and copying the files, or by performing a Copyback operation. For example:

    cp -a <*Snap2Dir*>/archivelog_dest*/ /oralogs/rman1/

    
    If you have previously chosen to replicate only one (not all) Oracle instance log locations, after the logs have been copied back, copy them to the other log location (if any) on the production server. For instance:

    If you have D:\ArchiveLogs1, F:\ArchiveLogs2, H:\ArchiveLogs3 and you initially selected to replicate only D:\ArchiveLogs1, after you copyback to D:\ArchiveLogs1, copy the same files to F:\ArchiveLogs2 and H:\ArchiveLogs3. You can either copy and paste them from D:\ArchiveLogs1 or run copyback again, specifying F:\ArchiveLogs2 and H:\ArchiveLogs3 as destinations, whichever is faster in your particular environment.
11. Set the ORACLE_SID and ORACLE_HOME environment variables.
12. Run the following commands as an Oracle user:

sqlplus "/ as sysdba"

- startup mount
- set autorecovery on;
- recover database until cancel using backup controlfile;
- alter database open resetlogs;

13. Database open will succeed after recovering the database.

Back to Top

Replication Set

Topics | How To | Related Topics

Overview

Configurable Properties

- Destination Host
- Data Replication Type
- Pre/Post Process
- Replication Options
- Orphan Files
- Filters
- User Security
- Advanced Options

Operations

- Manage Replication Pairs
- Change State
- Suspend Data Transfer
- Resume Data Transfer
- Manage Recovery Points
- Out Of Band Sync
- Monitor Data Replication

OVERVIEW

Replication Set represents a logical grouping of Replication Pairs. This allows you to configure certain properties and activities at the Replication Set level and thus control all the included Replication Pairs together.

Replication sets can be created in the following ways:

- Manually, see Create a Replication Set for step-by-step instructions.
- Using the Replication Set Creation Wizard. Create the Replication Set by right-clicking the ContinuousDataReplicator icon in the CommCell Browser and selecting **Replication Set Creation Wizard**.
- From a Replication Policy. A Replication Policy provides a sample template for the Replication Set. See Replication Policy for more details. See Create a Replication set from a Replication Policy for step-by-step instruction.

CONFIGURABLE PROPERTIES

The following are the properties which can be configured for a Replication Set.

DESTINATION HOST

Provides an option to Destination Host for all Replication Pairs within the Replication Set. Each Replication Set has one Destination Host. To replicate data to multiple hosts, you must create separate Replication Sets.

> When using an ONTAP filer as the destination you must specify a Windows User Account who is a member of the local Administrators group on the destination. See Create a Replication Set, for more information.

DATA REPLICATION TYPE

The following data replication types are available:

- **Continuous Replication -** See ContinuousDataReplicator for more information.
- **Disk Library Replication -** See Disk Library Replication for more information.
- **Scheduled Backup** - for Unix, to create replication sets for CDR to replicate data in Remote Backup mode.

PRE/POST PROCESS

You can specify commands to run either before a Recovery Point is created and/or after a Recovery Point is created. For general information about Pre/Post commands, see Pre/Post Processes.

USER ACCOUNTS

- You can define an account for accessing the destination directory.
- You can define an account with permissions to execute Pre/Post commands when creating Recovery Points.

See ContinuousDataReplicator in **User Accounts and Passwords** for more information.

REPLICATION OPTIONS

You can:

- Specify whether a mount point will automatically be created for the snapshots that comprise the latest Recovery Point each time one is created.
- Specify the maximum number of Recovery Points to be maintained at any time.
- Select the snap engine type for Recovery Point creation.
- Automatic backups of the Recovery Point with an appropriate Storage Policy and Backup Set.

DATA TRANSFER OPTIONS

A persistent connection is used as a data transfer mechanism across the network, to keep destination computer in sync with the defined content on the source computer. The following options for data transfer mechanisms are available:

- **Replication Compression -** allows you to compress the replicated data on the network. For more information, see Replication Compression.
- **Replication Encryption -** allows you to encrypt the replicated data on the network. For more information, see Replication Encryption.

ORPHAN FILES

Orphan files are files that are available only in destination directory and not in the source directory. Several options are provided to handle the orphan files. For more information, see Orphan Files.

FILTERS

Filters allow you to specify files, folders, and patterns to be excluded from data replication activities, for all Replication Pairs in a Replication Set. For more information, see Replication Filters.

USER SECURITY

Configure User Security for the Replication Set, after it has been created. For more information, see User Administration and Security.

ADVANCED OPTIONS

Configure the advanced properties for Replication Set to modify the default values of the minimum file size and the block size for hashing. For step-by-step instructions, see Create a Replication Set.

OPERATIONS

The following are the operations which can be initiated for a Replication Set.

MANAGE REPLICATION PAIRS

- You can create, delete, and edit Replication Pairs within the Replication Set.
- Specify the single destination computer that all Replication Pairs in the Replication Set will use. Note that different Replication Pairs can be configured for different destination paths within that same destination computer, as long as the source path of each Replication Pair is different.
- Manage the state of one or more Replication Pairs within the Set; start, suspend, resume, or abort data replication activities.
- Allows you to transfer the modified/new files on the source computer to the destination computer along with the data missing on destination. For more information, see Optimize Sync.

For more information, see Replication Pair.

CHANGE STATE

You can change the state of one or more Replication Pairs at once from the Replication Set level, see Change the state of Replication Pair for step-by-step

instructions.

SUSPEND DATA TRANSFER

- You can suspend Replication Pairs within a Replication Sets that have the same destination. See Suspend Data Replication Activity from the Replication Set for step-by-step instructions.

RESUME DATA TRANSFER

You can resume the Replication Pairs within a Replication Sets that have the same destination. See Resume Data Replication Activity for a Replication Set for step-by-step instructions..

MANAGE RECOVERY POINTS

- You can create a Recovery Point on-demand, or based on a schedule.
- You can browse Recovery Points. See Browse Data - ContinuousDataReplicator.

OUT OF BAND SYNC

It is possible to effect the initial transfer of data from a source to a destination without using the Baselining Phases. This can be useful when the connection between the source and the destination is constrained, such as a slow WAN connection. For step-by-step instructions to perform initial transfer of data, see Out Of Band Sync from the Replication Set.

MONITOR DATA REPLICATION

You can monitor all details of on-going replication activities for Replication Set, are shown in the Data Replication Monitor in the CommCell Console. For more information, see Monitoring Data Replication.

Replication Set - How To

Topics | How To | Related Topics

CREATE A REPLICATION SET

You can use the Wizard to create a Replication Set instead of using the procedure below; to do so, right-click the ContinuousDataReplicator icon in the CommCell, and select **All Tasks**, then select **Replication Set Creation Wizard**.

Required Capability: Capabilities and Permitted Actions

To create a Replication Set:

1. In the CommCell Browser, right-click the ContinuousDataReplicator icon of the source machine, select **All Tasks**, then select **Add Replication Set**.
2. In the General tab, specify a **Name** for the Replication Set, and a **Destination Host** that has CDR installed.

   On Windows, in the **User Name** field, specify a Windows User Account who has the following privileges:

   - Is a member of the local Administrators group on the destination.
   - Has read/write permissions to the product installation folders on the destination.
   - Has read/write permissions to the the replication logs directory on the destination if it is not in the default location.

- Has read/write permissions to the Job Results directory.

On Unix, in **Data Replication Type** area, select **Preserve Write Order** option for database application. This enables the use of full data journaling necessary to preserve the consistency of the database during replication.

For **Scheduled Replication** on Unix, the following options are selected by default to identify the backup data:

- **Use Journaling**: This option enables the use of internal journal to identify the backup data. Requires CDR driver; for supported platforms see System Requirements - ContinuousDataReplicator.
- **Create Snapshot**: This option enables the use of snapshots to identify the backup data. Requires a snapshot engine; for supported platforms see System Requirements - ContinuousDataReplicator - Snapshot Support.

For **Disk Library Replication**, select this option to replicate data between the MediaAgents installed in the source and destination computers.

3. In the Pre/Post Process tab, specify any commands to be run before a Recovery Point is created, and/or after one has been created, and specify a User Account with permissions to run the command(s) on the destination machine. For more information, see Pre/Post Processes.
4. In the Replication Options tab, make any appropriate selections. See the following for more information:
   - Enable or Disable Software Compression for a Replication Set
   - Configure the Replication Set for Data Encryption
   - Configure Recovery Points
   - Configure CDR for Backups of Recovery Points
5. Optionally, make any appropriate selections in the Storage Policy tab. For more information about Recovery Points and backups, see Recovery Points.
6. In the Orphan Files tab, specify how to treat data on the destination, which is not on the source machine. (See Data Replication - Orphan Files.)
7. Optionally, select the minimum file size to perform hash comparisons, and the block size for hashing in the Advanced tab.
8. Optionally, filter the content of any Replication Pairs that are created for this Replication Set in the Filters tab. For more information about CDR Filters, see Replication Filters. To exclude data from data replication operations, do one of the following:
   - Click the **Add** button and, in the **Enter Path** window, type the complete path (including drive letter) of the file/folder/directory that you want to exclude from data replication. Repeat this step if you want to add more files/folders/directories to the filter.
   - Click the **Browse** button and expand the file system of the client computer. Click the file/folder/directory that you want to exclude from data replication and then click **Add**. Repeat this step for each additional entry.
9. To configure User Security for the Replication Set, after it has been created, expand the ContinuousDataReplicator icon of the source machine, right-click the Replication Set and select **Properties**, then select the Security tab. For more information, see User Administration and Security.
10. Click **OK** to save your changes.

- To specify that all Replication Pairs created for this Replication Set are to use a Common Base Folder (the same folder on the destination host to contain all destination paths) see Edit a Replication Pair.

CREATE A REPLICATION SET FROM A REPLICATION POLICY

Before You Begin

- Review Replication Policy.

To create a Replication Set from a Replication Policy:

1. In the CommCell Browser, right-click the Agent icon of the source machine, select **All Tasks**, then select **Add Replication Set From Policy**.
2. In the **Select a Policy** dialog box, select the replication policy based on which you wish to create the replication set. The replication policies created for the corresponding agent will be displayed. Click **OK**.
3. In the **Enter Name** dialog box, specify the name of the replication set, and click **OK**.
4. The replication set will be created with the configurations available in the selected replication policy.

DELETE A REPLICATION SET

When you delete a Replication Set, all of the following are deleted as well:

- all its Replication Pairs
- all its Recovery Points
- all its mounted and shared snapshots

Required Capability: Capabilities and Permitted Actions

To delete a Replication Set:

1. Ensure that all Replication Pairs for the Replication Set you are deleting are in the "Stopped" or "New" state. If necessary, refer to Start/Suspend/Resume/Abort Data Replication Activity for instructions.
2. In the CommCell Browser, right-click a Replication Set icon, and select **All Tasks**, then select **Delete**.

   If Recovery Points were created for this Replication Set and any of the snapshots were mounted, select the **Unmount and delete any manually-mounted snapshots before deleting replication set** option.
3. Click **OK** in the confirmation screen.

START/SUSPEND/RESUME/ABORT DATA REPLICATION ACTIVITY FOR A REPLICATION PAIR

Before You Begin

- Review Monitoring Data Replication

Required Capability: Capabilities and Permitted Actions

To Start/Suspend/Resume/Abort a Replication Pair:

- In the CommCell Browser, right-click a Replication Pair, and select **Start, Start Full Resync, Suspend**, **Resume**, or **Abort**.

  From the Replication Pair level, you cannot Suspend a Replication Pair that is in the Replicating state; instead, do this from the Replication Set level. See, Suspend Data Replication Activity from the Replication Set for step-by-step instructions.

SUSPEND DATA REPLICATION ACTIVITY FROM THE REPLICATION SET

Before You Begin

- Review Monitoring Data Replication

Required Capability: Capabilities and Permitted Actions

To suspend all data replication activity from the Replication Set level:

1. In the CommCell Browser, right-click a **Replication Set** in the source machine, and select **All Tasks**, then select **Suspend Data Transfer**.
2. For Windows, if you **Suspend Data Transfer** for one Replication set, you will be prompted with a list of Replication Set that have the same destination.

   If you click **Yes** all Replication Pairs on that source computer which use the same destination computer will be placed in the Paused state, even those that are in different Replication Sets.

   For UNIX, data replication activities are suspended for all Replication Pairs in the Replication Set.

RESUME DATA REPLICATION ACTIVITY FROM THE REPLICATION SET

Before You Begin

- Review Monitoring Data Replication

Required Capability: Capabilities and Permitted Actions

To resume all data replication activity from the Replication Set level:

1. In the CommCell Browser, right-click a **Replication Set** in the source machine, and select **All Tasks**, then select **Resume Data Transfer**.
2. For Windows, you will be prompted with a list of Replication Set that have the same destination.

   Click **Yes** to resume all Replication Pairs between this source and destination - including the replication pairs in other replication sets.

   Click **No** to resume only the Replication Pairs in this replication set.

   For UNIX, data replication activities are resumed for all Replication Pairs in the Replication Set.

CONFIGURE THE REPLICATION SET FOR DATA ENCRYPTION

Before You Begin

- Encryption settings made at the Replication Set level are for encryption of data between the source machine and the destination machine.
- Encryption must be enabled at the client level prior to configuring data encryption for a Replication Set residing on that client. See Configure the Client for Data Encryption.

Required Capability: Capabilities and Permitted Actions

To configure data encryption for a Replication Set:

1. From the CommCell Browser, right-click the Replication Set and select **Properties**.
2. From the Replication Set Properties (Replication Options) tab, either select or clear **Encrypt During Data Transfer**.
3. Click **OK** to save your settings and close the Replication Set Properties.

Enable or Disable Software Compression for a Replication Set

Before you Begin

- Compression settings made at the Replication Set level are for compression of data between the source computer and the destination computer.

Required Capability: Capabilities and Permitted Actions

To enable/disable software compression for a Replication Set:

1. From the CommCell Browser, right-click the Replication Set and select **Properties**.
2. From the Replication Set Properties (Replication Options) tab, either select or clear **Compression ON**.
3. Click **OK** to save your settings and close the Replication Set Properties.

Add/Edit/Delete Filters for a Replication Set

Before You Begin

- When Filtering is edited for a Replication Set, any existing and active Replication Pairs will be automatically aborted and restarted to effect the change. If the newly configured filtering excludes files or folders that had previously been replicated, these files and folders on the destination(s) will be handled according to your Orphan Files settings.
- To configure filters for a new Replication Set, see Create a Replication Set.

Required Capability: Capabilities and Permitted Actions

To add/edit/delete a filter for a Replication Set:

1. In the CommCell Browser, right-click a Replication Set on the source machine, and select **Properties**.
2. In the Filters tab, specify the files and folders to filter from the content of any Replication Pairs that are created for this Replication Set. For more information, see Replication Filters.
3. Click **OK** to save your changes.

Specify a Snap Engine

Before You Begin

- Review Snapshot Engines - Support
- In the source computer:
  - For Windows 2003 or higher, VSS is used as the default snap engine.
- This procedure only applies to the destination computer.

Required Capability: Capabilities and Permitted Actions

To specify a snap engine on Windows to the destination computer:

1. In the CommCell Browser, right-click the Replication Set on the source computer and select **Properties**.
2. In the Replication Options tab, select one of the choices in the **Select Snap Engine Type for Recovery Point Creation** section.

   For Windows, **VSS** is the default selection; **ONTAP** can be used if available on the destination computer where the Recovery Point will be created.

   For UNIX, Snap Engine Type is automatically detected by the File System configured on the destination computer.
3. Click **OK** to save your changes.

On UNIX, CDR software automatically recognizes the File Systems configured on the destination computer and detects the appropriate snap engine. For more information on supported snap engines, see Snapshot Engines - Support.

Back To Top

Replication Pair - Windows

Windows | Unix | How To

OVERVIEW

A Replication Pair is used to map a designated set of data on a source computer to a location on a destination computer. This monitors change on the source and replicate the changes to the destination computer.

MANAGE REPLICATION PAIR

The activity of Replication Pairs can be controlled individually, or when acting from the Replication Set level, multiple Replication Pairs are controlled at the same time. Activities related to a Replication Pair include:

- Adding or Deleting a Replication Pair
- Changing the state of a Replication Pair; Start, Suspend, Resume, or Abort data replication
- Filtering content for all Replication Pairs in a Replication Set; see Replication Filters.

For step-by-step instructions, see Add a Replication Pair

COMMON BASE FOLDER

You can optionally specify a Common Base Folder, so that all destination paths for all Replication Pairs in the Replication Set will be created in the same folder on the destination host. This can be useful when you are adding multiple Replication Pairs at the same time, as in the case of replicating application data. For example:

If you specify `C:\replication\dest` as the base directory, your Replication Pairs might use the following as their destination directories:

```
C:\replication\dest\source_path_1\
C:\replication\dest\source_path_2\
C:\replication\dest\source_path_3\
```

If the `source_path_1` is specified as `C:\rep_pair_1` then the Replication Pair will use `C:\replication\dest\C\rep_pair_1` as its destination directory.

- If you change the settings for a Common Base Folder after any Replication Pairs have been created, the changes you make will affect only Replication Pairs created after that point; existing Replication Pairs will not be impacted by the change.
- A Common Base Folder is specified in the Manage Pairs dialog. For instructions about how to specify a Common Base Folder, see Add a Replication Pair.
- For more information about replicating application data, see Application Integration.

REPLICATION PAIR CONFLICTS

Certain Replication Pair combinations are not possible, or are only possible using different Replication Sets. The following table highlights the most common such cases, without enabling chaining for replication option in a control panel:

TYPE OF COMBINATION OR CONFLICT	SAME REPLICATION SET	DIFFERENT REPLICATION SET	EXAMPLES NOTE THAT (A), (B), (C) IN THE EXAMPLES REPRESENT DIFFERENT COMPUTERS	NOTES
Exact same destination path	No	No	(A) `D:\test` => (B) `E:\test` (A) `F:\test` => (B) `E:\test`	
Child/parent destination paths	No	No	(A) `D:\test` => (B) `E:\test` (C) `F:\test` => (B) `E:\test\test1` - - - - - - (A) `D:\test` => (B) `E:\test` (C) F:\ => (B) E:\	

Source path same as destination (volumes are on different computers)	N/A	Yes	(A) D:\test => (B) D:\test (C) D:\test => (B) E:\test	
Same destination path on multiple computers	N/A	Yes	(A) D:\test => (B) D:\test (A) D:\test => (C) D:\test	
Same destination volume, different folders	Yes	Yes	(A) D:\test => (B) D:\test (C) E:\data => (B) D:\test1	
Same source path in multiple pairs	No	Yes	(A) D:\test => (B) D:\test (A) D:\test => (C) E:\test	
Source parent folder already used (Nested mount point is not supported.)	Yes	Yes	Existing: (A) D:\ => (B) D:\ New: (A) D:\data => (B) E:\data	The new pair will change to the parent path.
Source child folder already used	No	No	Existing: (A) D:\data => (B) E:\data New: (A) D:\ => (B) D:\	
Destination parent folder already used	No	No	Existing: (A) E:\src => (B) G:\des New: (B) G:\des => (C) I:\ (B) G:\ => (C) I:\des1 (B) G:\data => (C) H:\test (B) G:\des\test => (C) I:	

You can enable chaining for replication to use destination paths as a source for replication pairs. For step-by-step instructions to enable chaining, see Enable Chaining for Replication.

For chaining, replication pair source and destination paths can be configured as follows:

(A) E:\test => (B) G:\test

(B) G:\test => (C) H:\data

(C) H:\data => (D) I:\test

Ensure that your destination data is not replicated back to any volume that is being used as a source as this may cause an endless loop. For example:

(A) E:\test => (B) G:\test

(B) G:\test => (C) H:\data

(C) H:\data => (D) I:\test

(D) I:\test => (A) E:\test or E:\

(OR)

(A) D:\test => (B) G:\data

(A) F:\test1 => (B) H:\data1

Note that when you create a new replication pair(s) as below then it will from an endless loop.

(B) G:\data => (A) F:\test1

(B) H:\data1 => (A) D:\test

Here (A), (B), (C), (D) are different computers

IMPORTANT CONSIDERATIONS

When creating a Replication Pair, and selecting the source and destination paths for replication, consider the following:

Attempting to use CDR to protect the system volume of a Windows computer is strongly discouraged. Due to the normally high rate of I/O on such a volume (e.g., the C: drive), it would prove very difficult to create a snapshot of it, since a sufficient period of no disk activity is required in order to create the snapshot. Similarly, a system volume is an inappropriate choice as a destination for data replication, and is also strongly discouraged.

SOURCE

- A specific directory or volume on a source computer can be replicated multiple times to different destination paths on either the same/different destination computer, or different destination computers, by creating a separate Replication Set and Pair for each different destination path.

  For instance, Replication Set 1 could have a Replication Pair configured as follows:

  - Source on Computer A: D:\
  - Destination on Computer B: G:\A_to_B_1\

Replication Set 2 could have a Replication Pair configured as follows:

- Source on Computer A: `D:\`
- Destination on Computer B: `H:\A_to_B_2\`

- When source data contains stubs from a data migration, the stubs themselves will be replicated, and not the actual data which has been migrated.
- For a Replication Pair with the **Automatically detect mount points** option selected, for each mount point in the specified directory or volume, an additional, separate Replication Pair will be automatically created and its source path marked with an asterisk ("*") at the beginning to indicate it is a mount point.
- It is possible to effect the initial transfer of data from a source to a destination without using the Baselining Phases. This can be useful when the connection between the source and the destination is constrained, such as a slow WAN connection. For step-by-step instructions to perform initial transfer of data, see Out Of Band Sync from a Replication Pair.

DESTINATION

- The source path cannot be a mapped network drive.
- On Windows, a UNC path can be used as a destination. For data replication to succeed, the destination computer must have permissions to the location specified by the UNC path. Note that the UNC path can be specified to a NetApp file server. For more information, see ONTAP Snapshot for ContinuousDataReplicator.
- When selecting a destination folder for a Replication Pair, ensure that the selected path does not contain any data as the existing data will be pruned during Orphan File processing.
- It is recommended that, you do not include UNC path and iSCSI LUNs as destination path in the same Replication Set.

Back to Top

Replication Pair - How To

Windows | **Unix** | **How To**

Add a Replication Pair

Edit a Replication Pair

Delete a Replication Pair

Enable Chaining for Replication

ADD A REPLICATION PAIR

Required Capability: Capabilities and Permitted Actions

To add a Replication Pair:

1. In the CommCell Browser, right-click a Replication Set in the source machine, and select **All Tasks**, then select **Manage Pairs**.
2. To add a Replication Pair, from the Manage Pairs dialog box, perform the following:
   - Specify a Common Base Folder so that all destination paths for all Replication Pairs in the Replication Set will be created in the same location on the destination host.
   - For File System data, click **Add**. In the Add Replication Pair dialog box, type or click **Browse** to navigate to and select, a source volume or directory for the Source Path fields, then click **OK**.
   - For application data, click **Add App.** In the Add Applications dialog box, select all of the directories associated with the application data to be replicated, and click **OK**. For more information about replicating application data, see Application Integration.
3. To specify that additional Replication Pairs are to be created for mount points that are part of your selected source volume or directory, select **Automatically detect mount points**. Mount points will be ignored if this option is not selected.
4. Select **Include files that do not match with destination copy** option to specify the files that already exist on the destination and transfer the new/modified files along with data missing on destination.
5. When your changes are complete, click **Apply**.

Filtering of content for Replication Pairs is configured at the Replication Set level; see Replication Filters.

EDIT A REPLICATION PAIR

Before You Begin

- You can edit a Replication Pair which is Stopped or Failed. See Monitoring Data Replication and Start/Suspend/Resume/Abort Data Replication Activity.

Required Capability: Capabilities and Permitted Actions

To edit a Replication Pair:

1. In the CommCell Browser, right-click a Replication Set in the source machine, and select **All Tasks**, then select **Manage Pairs**.
2. In the **Manage Pairs** dialog, select the pair to edit, click **Edit**.
3. Make the necessary changes in the **Edit Replication Pair** dialog box.
4. When your changes are complete, click **Apply**.

DELETE A REPLICATION PAIR

Before You Begin

- You can delete the Replication Pair when it is Stopped. See Monitoring Data Replication and Start/Suspend/Resume/Abort Data Replication Activity.
- When you delete a Replication Pair, its Recovery Points are not deleted. See Delete a Recovery Point.

Required Capability: Capabilities and Permitted Actions

To delete a Replication Pair:

Select one of the following methods:

1. From the CommCell Browser, right-click a Replication Pair, and select **Delete**.
2. From the CommCell Browser, right-click a Replication Set in the source machine, and select **All Tasks**, then select **Manage Pairs**.

   In the **Manage Pairs** screen, select a Replication Pair, and click **Delete**.

   When your changes are complete, click **Close**.

   After deletion of individual Replication Pair in a Replication Set, you have to manually select the **Include files that do not match with destination copy** option for an existing Replication Pair. To enable **Include files that do not match with destination copy** option, see Edit a Replication Pair.

ENABLE CHAINING FOR REPLICATION

To enable Replication Chaining:

1. From the CommCell Browser, right-click the CommServe icon, click **Control Panel**, and then click **Replication and Workstation Settings**. This opens the Replication and Workstation Settings.
2. Select the **Advanced** tab.
3. Check **Enable chaining for Replication and WorkStation Backup clients** option.
4. Click **OK**.

Back To Top

Application Integration

OVERVIEW

CDR will assist in configuring Replication Pair content, by automatically discovering directories that should be replicated (logs, databases, etc.) to protect supported applications. Note that all replication of application data is performed while the server is online. (See Supported Data Types.) Consistent Recovery Points, which define a point-in-time where application data is in a consistent state, can be created using CDR, ensuring the application data can be restored to that point-in-time. After data for these applications has been replicated, when browsing the data from the associated application *i*DataAgent (e.g., browsing from the Exchange *i*DataAgent), any Consistent Recovery Points maintained by CDR will be displayed along with traditional backups and/or any QR Volumes that may be available for that application's data.

When replicating application data, multiple Replication Pairs are created at the same time. It can be useful to specify a Common Base Folder, so that all destination paths for all Replication Pairs in the Replication Set will be created in the same location on the destination host. For more information, see Common Base Folder.

NON-INTEGRATED APPLICATIONS

For applications that are not integrated with CDR, their data can still be replicated by appropriately configuring Replication Sets and Pairs to capture all of the associated data. For more information, see Consistent Recovery Points.

OFFLINE MINING

CDR can also replicate application data for offline mining of Active Directory, Exchange, and SharePoint data. The steps involved in replicating application data for offline mining are the same as those described in Replicating Application Data using ContinuousDataReplicator; however, certain configuration considerations apply depending on the application being replicated. Review the following offline mining overview topics for more information on replicating application data for offline mining in your environment:

- Active Directory Offline Mining
- Exchange Offline Mining Tool
- SharePoint Offline Mining

REPLICATING APPLICATION DATA USING CONTINUOUSDATAREPLICATOR

Refer to the following for step-by-step instructions for replicating application data:

- Configure CDR to Replicate Application Data
- Change Account for Accessing Application Servers/Filers
- Using ContinuousDataReplicator with Microsoft Exchange
- Using ContinuousDataReplicator with Microsoft SQL Server
- Using ContinuousDataReplicator with Oracle

IMPORTANT CONSIDERATIONS

GENERAL

- It is recommended that you do not change the name of the client during installation as this can cause the application discovery to fail.

SQL DATA

- When you add or edit a Replication Pair, and you select **Add App** in the Manage Pairs dialog, SQL volumes will not be discovered if any databases are in the *suspect* or *loading* state. Use the SQL Enterprise Manager to change the mode of the database(s).
- During **Add App,** SQL Backward compatibility Tool should be installed in order to detect SQL databases created using higher version for SQL.
- **Add App** discovers the location of system databases on the client, as well as user-defined databases (`.mdf`, `.ndf`, `.ldf`), which you can select for data replication.
- The recommended maximum number of SQL databases in an individual replication set should not exceed 150.

ORACLE DATA

- If there is more than one log file destination, CDR will discover all of them, and create all the necessary Replication Pairs.
- Ensure that all `ARCHIVE_LOG_DEST_n` for the Oracle database are set to a path, not a parameter, or the Oracle discovery may fail.

  Correct example, set to a path:

  ```
  ALTER SYSTEM SET LOG_ARCHIVE_DEST_10 = 'LOCATION=F:\backup2\logs';
  ```

  Incorrect example, set to a parameter:

  ```
  ALTER SYSTEM SET LOG_ARCHIVE_DEST_10 = USE_DB_RECOVERY_FILE_DEST;
  ```

- Oracle volumes configured with Symantec Storage Foundation QIO/ODM software are not supported for CDR on Unix.

Back to Top

Application Integration - How To

Topics | How To | Support

CONFIGURE CDR TO REPLICATE APPLICATION DATA

Before You Begin

- Review Application Integration

Required Capability: Capabilities and Permitted Actions

To configure CDR to replicate application data:

1. In the CommCell Browser, right-click the ContinuousDataReplicator icon of the source machine, and select **Properties**.
2. Select the **Authentication** tab of the **Agent Properties** screen, and select one of the following:
   - **Exchange/SQL** tab - select either **Microsoft Exchange Server** or **Microsoft SQL Server**, click **Edit**, then supply the necessary credentials for the server. For more information, see specify the Account for Accessing Application Servers/Filers.

     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.
   - **Oracle** tab - click **Add Instance** and supply the necessary credentials and information for Oracle. Supply a user account with administrator privileges to access the Oracle application.

     For Oracle on Unix, you can use the operating system user account to verify the rights to perform all data protection and recovery operations for the associated Oracle instance.

     For Oracle on Windows, you can click the **Change** button and use the Impersonate User option for this purpose.
3. Click **OK** to save your changes.

To create a Replication Pair for application data, see Add a Replication Pair.

CHANGE ACCOUNT FOR ACCESSING APPLICATION SERVERS/FILERS

Required Capability: See Capabilities and Permitted Actions

To change most Agent user account for accessing an application server or filer:

1. From the CommCell Browser, expand the tree if necessary to view the affected agent icon. Then right-click the agent icon and click **Properties**.
2. From the General tab, click the **Change Account** button associated with the User Account information.
3. For Active Directory, if you want to use NTLM Bind, Enter **Administrator** as the user and then enter the administrator's password in the Change User Account dialog box.
4. For NetWare/NDS, File Archiver for NetWare, or Active Directory, enter the requested information in the Change User Account dialog box.
5. For the affected Exchange-based agent or QR Agent, enter the requested information in the Exchanged-based Change User Account dialog box.
6. For SharePoint, enter the requested information in the SharePoint-based Change User Account dialog box.
7. Click **OK** to save the settings.

To change a Quick Recovery Agent user account for accessing an application server or filer:

1. From the CommCell Browser, expand the tree if necessary to view the affected agent icon. Then right-click the agent icon and click **Properties**.
2. Click the Authentication tab.
3. From the default **Exchange/SQL** sub-tab, click the appropriate application and then click **Edit**.

   When configuring the Exchange server(s) for the Quick Recovery Agent on a cluster, be sure to enter the Exchange server name into the **Exchange Server Name** field of the **Change User Account** dialog box. If you do not enter the server name, the agent may not be able to detect the Exchange Server.
4. Enter the required information in the Change User Account dialog box.
5. Click **OK**.

To change a ContinuousDataReplicator user account (Windows only) for accessing an application server:

1. From the CommCell Browser, expand the tree if necessary to view the affected agent icon. Then right-click the agent icon and click **Properties**.
2. Click the Authentication tab.
3. Click the appropriate application in the list and then click **Edit**.
4. Enter the required information in the Change User Account dialog box.
5. Click **OK**.

USING CONTINUOUSDATAREPLICATOR WITH MICROSOFT EXCHANGE

The following section provides the steps required to use CDR for data replication and recovery for Microsoft Exchange data based on a single source and single destination. If your environment uses a different scenario, adjust your steps accordingly.

BEFORE YOU BEGIN

- A full understanding of the following subjects will prove helpful:
  - Overview - ContinuousDataReplicator
  - Data Replication
  - Application Integration
  - Recovery Points
  - Recover Replicated Data

DATA REPLICATION AND CONSISTENT RECOVERY POINTS FOR MICROSOFT EXCHANGE

Required Capability: Capabilities and Permitted Actions

To use CDR to replicate Microsoft Exchange data and create Consistent Recovery Points:

1. Select two computers on which to install CDR, one designated as the source computer, and one designated as the destination computer.
   - Verify that they both meet the System Requirements.
   - Install the ContinuousDataReplicator software for Windows on both computers.
2. If you are using VSS, you may want to use a separate shadow storage area for Consistent Recovery Points since exchange uses a large amount of cache space.
3. If you are replicating data for a clustered Exchange Server, and you are using the legacy offline method to quiesce the Exchange Server to facilitate Consistent Recovery Point creation, CDR will need to be configured with account information for the Exchange Server; for instructions, see Change Account for Accessing Application Servers. If you are using VSS to quiesce the Exchange Server, you can ignore this step.
4. Configure CDR to Replicate Application Data.
5. Create a Replication Set.
6. Configure Consistent Recovery Points for the Microsoft Exchange data; see Configure CDR Recovery Points.
7. Add a Replication Pair.
8. Start Data Replication Activity.
9. Create a Consistent Recovery Point.

Additional Recommendations

- On both the source and destination computers, it is recommended that you Configure Throttling for CDR Replication Activities.
- It is recommended that you Configure Alerts. For more information, see Alerts.

  When using QSnap on a source computer it is recommended that you also see Space Check for the Quick Recovery and ContinuousDataReplicator Agents and configure the `Disk Space Low` alert to provide warning that the source computer is running out of disk space, which will ultimately cause replication activity to be System Aborted.
- Optionally, you can Configure CDR for Backups of Recovery Points.

  To perform backups of Recovery Points, you must also install the Windows File System iDataAgent on both the source and destination computers.
- Monitor Data Replication Activities.
- View the Recovery Point Creation History.
- Run and review the CDR Replication Job Summary Report.

COPYBACK OF MICROSOFT EXCHANGE DATA

Exchange data is restored at the Storage Group level. While you can restore Exchange data from a Recovery Point, a backup of a Recovery Point, or the Live Copy, these methods will not ensure consistency of the application data; only a restore from a Consistent Recovery Point, or a backup of one, will ensure consistency of application data. Copyback is recommended as the primary method of moving the replicated data back to the production Exchange Server, in addition to restoring a backup of a Consistent Recovery Point where that is appropriate. Exchange circular logging can be either enabled or disabled, with no impact on the Copyback operation. To ensure application integrity, you must use **Add App** to create your Replication Pairs. (Refer to Add or Edit a Replication Pair and Application Integration.) **Add App** discovers the location of the Exchange `.chk` file and `tmp.edb` file, which means the Exchange `.chk` file will not have to be deleted before performing the Copyback operation, since it is restores from the same point-in-time as the database and log files.

For step-by-step instructions, see Copy Back Exchange Data from a Consistent Recovery Point.

Additional Recommendations

- View the Recovery Point Copyback History.
- Run and review the CDR Copyback Job Summary Report.

USING CONTINUOUSDATAREPLICATOR WITH MICROSOFT SQL SERVER

The following section provides the steps required to use CDR for data replication and recovery of Microsoft SQL Server data based on a single source and single destination. If your environment uses a different scenario, adjust your steps accordingly.

BEFORE YOU BEGIN

- A full understanding of the following subjects will prove helpful:
  - Overview - ContinuousDataReplicator
  - Data Replication
  - Application Integration
  - Recovery Points

- Recover Replicated Data

DATA REPLICATION AND CONSISTENT RECOVERY POINTS FOR MICROSOFT SQL SERVER

Required Capability: Capabilities and Permitted Actions

To use CDR to replicate Microsoft SQL Server data and create Consistent Recovery Points:

1. Select two computers on which to install CDR, one designated as the source computer, and one designated as the destination computer.
   - Verify that they both meet the System Requirements.
   - Install the ContinuousDataReplicator software for Windows on both computers.
2. For VSS, you may want to use separate volume for shadow storage.
3. Configure CDR to Replicate Application Data.
4. Create a Replication Set.
5. Configure Consistent Recovery Points for the SQL Server data; see Configure CDR Recovery Points.
6. Add a Replication Pair.
7. Start Data Replication Activity.
8. Create a Consistent Recovery Point.

Additional Recommendations

- On both the source and destination computers, it is recommended that you Configure Throttling for CDR Replication Activities.
- It is recommended that you Configure Alerts. For more information, see Alerts.

  It is also recommended that you also see Space Check for the Quick Recovery and ContinuousDataReplicator Agents and configure the `Disk Space Low` alert to provide warning that the source computer is running out of disk space, which will ultimately cause replication activity to be System Aborted.
- Optionally, you can Configure CDR for Backups of Recovery Points.

  To perform backups of Recovery Points, you must also install the Windows File System iDataAgent on both the source and destination computers.
- Monitor Data Replication Activities.
- View the Recovery Point Creation History.
- Run and review the CDR Replication Job Summary Report.

COPYBACK OF MICROSOFT SQL SERVER DATA

SQL data is restored at the database level. While you can restore SQL data from a Recovery Point, a backup of a Recovery Point, or the Live Copy, these methods will not ensure consistency of the application data; only a restore from a Consistent Recovery Point, or a backup of one, will ensure consistency of application data. Copyback is recommended as the primary method of moving the replicated data back to the production SQL Server, in addition to restoring a backup of a Consistent Recovery Point where that is appropriate. To ensure application integrity, you must use **Add App** to create your Replication Pairs. (Refer to Add or Edit a Replication Pair and Application Integration.) **Add App** discovers the location of, not only user-defined databases (`.mdf`, `.ndf`, `.ldf`) but also any system databases on the client, which you can select for data replication.

For step-by-step instructions, see Copy Back SQL Data from a Consistent Recovery Point.

Additional Recommendations

- View the Recovery Point Copyback History.
- Run and review the CDR Copyback Job Summary Report.

USING CONTINUOUSDATAREPLICATOR WITH ORACLE

The following section provides the steps required to use CDR for data replication and recovery of Oracle data based on a single source and single destination. If your environment uses a different scenario, adjust your steps accordingly.

BEFORE YOU BEGIN

- A full understanding of the following subjects will prove helpful:
  - Overview - ContinuousDataReplicator
  - Data Replication
  - Application Integration
  - Recovery Points
  - Recover Replicated Data

DATA REPLICATION AND CONSISTENT RECOVERY POINTS FOR ORACLE

Required Capability: Capabilities and Permitted Actions

To use CDR to replicate Oracle data and create Consistent Recovery Points:

1. Select two computers on which to install CDR, one designated as the source computer, and one designated as the destination computer.
   - Verify that they both meet the System Requirements.
   - Install the ContinuousDataReplicator software on both computers.
2. Consistent Recovery Points can only be created for Oracle if the database and archive logs are not on the same volume on the source computer (the production server.) The archive logs can be in more than one location, as long as none of them is on the same volume as the database. If necessary, use the Oracle `ALTER` command to set the archive log location to a different volume. For example:

   ```
   ALTER SYSTEM SET LOG_ARCHIVE_DEST_10 = 'LOCATION=F:\ORALOG1';
   ```
3. If you are using QSnap, consider the following:
   - To use QSnap with CDR on UNIX, before you can begin creating Replication Sets and Replication Pairs, you must first configure source and destination volumes as CXBF devices. For more information, see QSnap for ContinuousDataReplicator.
   - For VSS, you may want to use separate volume for shadow storage.
4. Configure CDR to Replicate Application Data.
5. Create a Replication Set.
6. Configure Consistent Recovery Points for the Oracle data; see Configure CDR Recovery Points.
7. Add a Replication Pair. The database itself, as well as all Oracle log destinations, e.g., `archivelog_dest1`, `archivelog_dest2`, etc., must be replicated to ensure that a Copyback operation can succeed. When you use the **Add App** button during Replication Pair creation, all log file location will be discovered, and all necessary Replication Pairs will be created.
8. Start Data Replication Activity.
9. Create a Consistent Recovery Point. Each Consistent Recovery Point for a given Oracle database will comprise two snapshots:
   - The first snapshot will include all Oracle data volumes.
   - The second snapshot will include all Oracle log volumes and backup control file for the database at the time the Consistent Recovery Point was created.
   - If Replications Pairs for Oracle log and database replication are configured to use the same destination volume, the actual number of Recovery Points retained will be the number specified in **Maximum Number of Recovery Points** divided by two, because separate snapshots will be created for the logs and the database. For example, if you have specified 10 as the maximum number of Recovery Points, but use the same destination volume for the logs and database, only 5 Recovery Points will be retained, with 2 snapshots in each.

Additional Recommendations

- On both the source and destination computers, it is recommended that you Configure Throttling for CDR Replication Activities.
- It is recommended that you Configure Alerts. For more information, see Alerts.

  It is also recommended that you also see Space Check for the Quick Recovery and ContinuousDataReplicator Agents and configure the `Disk Space Low` alert to provide warning that the source computer is running out of disk space, which will ultimately cause replication activity to be System Aborted.
- It is recommended that you do not configure Oracle log and database to use the same destination volume.
- Optionally, you can Configure CDR for Backups of Recovery Points.

  To perform backups of Recovery Points, you must also install either the Windows File System *i*DataAgent on both the source and destination computers, or the Unix File System *i*DataAgent on both the source and destination computers. You cannot replicate Windows data to a UNIX computer, nor the converse.
- Monitor Data Replication Activities.
- View the Recovery Point Creation History.
- Run and review the CDR Replication Job Summary Report.

COPYBACK OF ORACLE DATA

Oracle data is restored at the database level. While you can restore Oracle data from a Recovery Point, a backup of a Recovery Point, or the Live Copy, these methods will not ensure consistency of the application data; only a restore from a Consistent Recovery Point, or a backup of one, will ensure consistency of application data. A Consistent Recovery Point is recommended as the source when moving replicated data back to the production Oracle server. To ensure application integrity, you must use **Add App** to create your Replication Pairs. (Refer to Add or Edit a Replication Pair and Application Integration.) **Add App** discovers the location of all user-defined and system databases on the client, which you can select for data replication.

For step-by-step instructions, see Copy Back Oracle Data from a Consistent Recovery Point.

For Unix, sparse files attributes are not retained during Copyback, neither from a Live Copy nor a Recovery Point; the files assume the attributes of regular files on the recovery host.

Additional Recommendations

- View the Recovery Point Copyback History.
- Run and review the CDR Copyback Job Summary Report.

Back to Top

QSnap for ContinuousDataReplicator

Topics | Related Topics

Overview

- CDR on UNIX

QSnap for CDR on UNIX

- Install QSnap
- Configure UNIX CXBF Devices for CDR
- Deconfigure UNIX CXBF Devices used by CDR

OVERVIEW

QSnap can be used with ContinuousDataReplicator on a source computer or a destination computer for several purposes. There are differences in its requirement and use, depending on the operating system, which are explained in some detail below. For CDR or QSnap installation procedures, refer to Deployment.

CDR ON UNIX

QSnap can be optionally installed with CDR on Linux, on a destination computer, to perform the functions described below. In some cases this can provide better performance and functionality than using Logical Volume Management (LVM) to create File System Snapshots.

- On a destination computer, QSnap can be used to create the snapshots that comprise a Recovery Point. By default, LVM is be used for this function instead of QSnap; for more information about LVM, refer to your Linux documentation.

For CDR on AIX, QSnap is not supported for these functions; refer to IBM documentation for more information about File System Snapshots created by LVM.

For CDR on HP-UX, QSnap is not supported for these functions.

QSnap can be installed with CDR on Solaris, on a destination computer, for configuring UFS and VxFS as CXBF devices.

QSNAP FOR CDR ON UNIX

By default, CDR on UNIX utilizes Logical Volume Management (LVM) for snapshot purposes, for the creation of Recovery Points on the destination computer. QSnap can be used for these purposes instead of LVM; to do so, you must install QSnap and configure CXBF devices.

INSTALL QSNAP

In addition to the CDR software, the QSnap software must be installed as well, on the destination computer(s). This can be done at the same time you install CDR, or separately. For instructions, see Install QSnap - Unix.

To change which snap engine will be used, see Specify a Snap Engine for a CDR Source or Destination Computer.

CONFIGURE UNIX CXBF DEVICES FOR CDR

Before any volume can be used as a destination for CDR with QSnap on UNIX, it must be mounted and configured as a CXBF device using the procedure detailed here.

1. Ensure that CDR and QSnap have been installed on the computer. For procedures, see Install ContinuousDataReplicator - UNIX and Install QSnap - Unix.
2. Mount any volumes that will be used as a destination.
3. Use Volume Explorer to configure CXBF devices for ContinuousDataReplicator.

   For CDR on Solaris, you have the option to configure CXBF devices during the creation of Replication Sets or Replication Pairs. This option applies to UFS and VxFS.

4. Once you have configured destination volumes as CXBF devices, you can create Replication Sets and Replication Pairs that utilize those volumes.

   When selecting destination volumes for Replication Pairs, you must ensure that each volume you select is a CXBF device.

DECONFIGURE UNIX CXBF DEVICES USED BY CDR

If you no longer want to use a volume as a CXBF device for CDR, you can deconfigure it using Volume Explorer.

Back to Top

VSS for ContinuousDataReplicator

Topics | Related Topics

OVERVIEW

VSS is utilized on the source computer for creating snapshots and the VSS writers are used for performing quiesce/unquiesce of Exchange and SQL data. VSS is also used for creating Recovery Points. For Windows 2000 or Windows XP, snapshots on the source computer during the SmartSync Scan phase are created by QSnap, and the default quiesce method is used for quiescing the applications. Recovery Point snapshots on the destination computer can be created using VSS to create Shadow Copies for the snapshots.

If the source computer has Calypso VSS Provider installed, then the VSS quiesce method is used to perform quiescing the applications. For Exchange and SQL Server's, Calypso VSS Provider is utilized on the source computer to perform an online quiesce/unquiesce, instead of using CDR's default offline quiesce/unquiesce, while a marker is placed in the log file for a Consistent Recovery Point. Online quiesce allows the server to continue to function with no interruption for users. If CDR VSS Provider is not installed then the legacy quiesce method is used for creating a Consistent Recovery Point.

CONFIGURATION

The following is necessary to implement this functionality:

1. Verify that the source or destination computer meets the System Requirements, and is a supported operating system platform and application server software for using VSS with CDR. Refer to ContinuousDataReplicator - Application Support.
2. Install CDR software with the **VSS Provider for CDR** on the source and/or destination computer(s). Also, install CDR software on the destination computer; the **VSS Provider for CDR** is not required on the destination. For instructions, see Deployment - ContinuousDataReplicator.
3. For a destination computer, to use VSS to create Shadow Copies for Recovery Points, right-click the Replication Set on the source computer and select **Properties**. In the Replication Options tab, select one of the choices in the **Select Snap Engine Type for Recovery Point Creation** section.

   If you have the hardware VSS provider setup as the default snap engine and you want to use Microsoft VSS for creating snapshots on the destination computer, you can use the `nUseVSSSoftwareProvider` registry key. Note that the CDR services will need to be recycled on the client computer for changes to take affect.

VSS CONSIDERATIONS

Before using VSS, review the following information:

- VSS is the default snap engine used on a Windows 2003 or higher source computer.
- VSS is not available with Windows 2000. Do not specify VSS as the Snap Engine for a computer running Windows 2000.
- The CDR VSS Provider is supported on a cluster and only needs to be installed on the physical node of a cluster.
- The VSS Provider for CDR does not support Multi Instancing; it must only be installed to Instance001.
- A VSS online quiesce ensures that users will not lose their connection to the server, e.g., Outlook clients will not be disconnected.
- Exchange can be quiesced at the Storage Group level.
- When creating Recovery Points using VSS, a maximum of of 256 snapshots are allowed.
- Recovery Point snapshots created with VSS can be deleted in any order.
- The VSS cache can be configured using the **`vssadmin add shadowstorage`** command from a command line prompt. Refer to Microsoft documentation for details.
- When using VSS, it is recommended that you do not start a large number of pairs, using the same source volume, simultaneously. To avoid any issues, you should vary the starting time for pairs which use the same source volume.

LICENSE REQUIREMENT

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.

The VSS Provider for CDR does not require a license.

Back to Top

File System Snapshot

Topics | How To | Related Topics

OVERVIEW

File System Snapshot, a built-in feature of the UNIX operating system, is supported by the Quick Recovery Agent, Recovery Director and ContinuousDataReplicator on UNIX, and is used to create point-in-time snapshots of volumes. These snapshots do not involve bitmaps, so only full volume copies are possible; you cannot create incremental snapshots.

For the QuickRecovery Agent on UNIX, the system copies a File System Snapshot to a secondary volume through the use of a Generic Enabler snapshot script that you create. Another Generic Enabler script enables you to recover the snapshot. If the snapshot is on the local machine, you can do a quick recovery. If you have copied the snapshot to a remote machine, you can recover using Copyback. The Generic Enabler feature and sample scripts are provided with the Quick Recovery Agent.

For CDR on UNIX, File System Snapshots are employed in the SmartSync Scan phase, and can also be used when creating Recovery Point on a destination computer. For more information about about the SmartSync Scan phase see Job Phases; for more information about Recovery Points, see Recovery Points.

- For the QuickRecovery Agent, to copy to a remote machine or copy back from a remote machine, both the local and the remote machine must have the QuickRecovery Agent installed.
- In Linux, File System Snapshots are managed by Logical Volume Manager (LVM); therefore, they will only work with a mounted logical volume and not with a cxbf device. If a cxbf device was mistakenly mounted on a logical volume, follow the procedures in Deconfigure a CXBF Device in Volume Explorer. If that fails to deconfigure the device, see Defunc and Delete a CXBF Device.

CONFIGURATION

To configure File System Snapshot for the QuickRecovery Agent on UNIX:

1. Install the QuickRecovery Agent on a supported platform.
2. Create a QR Policy; select `Generic snapshot on Unix` as the Snapshot Engine Type.
3. Configure the QuickRecovery Agent subclient to use your Generic snapshot script.

To configure File System Snapshots for CDR on UNIX:

1. Install ContinuousDataReplicator on a supported platform.
2. Specify a Snap Engine on both the source and destination, as needed.

LICENSING

A product license for either the QuickRecovery Agent or ContinuousDataReplicator is required on the source and destination computers; no extra license is required for the File System Snapshot feature.

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

ONTAP Snapshot for ContinuousDataReplicator

Topics | Related Topics

OVERVIEW

The ONTAP Snapshot is used to create Recovery Point snapshots of a destination. The destinations are UNC path, iSCSI LUNs to a NetApp file server in case of Windows and a mounted NFS export in case of Unix.

CONFIGURATION

The following must be done before using ONTAP Snapshot.

1. **Verify System Requirements** -- The source and destination computers must meet the System Requirements, and have a supported operating system platform and application server software for using ONTAP Snapshot with ContinuousDataReplicator (CDR). Refer to ContinuousDataReplicator -- Application Support and Snapshot Engines - Support.
2. **Install the Software** -- Install CDR software on the source and destination computer(s). For instructions, see Deployment - ContinuousDataReplicator.
3. **Configure the Agent** -- For Unix, right-click the Agent and select **Properties**. In the Authentication tab, click **Filer Management.** From the Array management dialog box, add the credentials for the user account to be used when using ONTAP snapshots.

   Ensure that you provide the fully qualified domain name or TCP/IP address of the filer in the **Control Device ID** field.

4. **Configure the Replication Set** -- For Windows, for a destination computer to use ONTAP Snapshot to create Recovery Points, right-click the Replication Set on the source computer and select **Properties**. In the Replication Options tab, select **ONTAP** in the **Select Snap Engine Type for Recovery Point Creation** section. To access the filer, specify user account information in the General tab of the replication set properties on the source computer. For Unix, the Snapshot engine is automatically selected based on your destination.

   For Fan-In Recovery Points on Windows, right-click the ContinuousDataReplicator Agent on the destination computer and select **Properties**. In the Fan-In tab, click **Recovery Points** tab and select **Use ONTAP snapshot for ONTAP LUN destinations** option, to create ONTAP snapshots on iSCSI LUN. If your destination is NetApp filer and destination path is UNC path, ONTAP Snapshot engine is automatically selected.

   To access the filer, you must specify the user credential details as follows:
   - For Recovery Points, from the General tab of the replication set properties on the source computer, specify the credentials for the user account to be used when using ONTAP snapshots.
   - For Fan-In Recovery Points, from the General tab of the agent properties on the destination computer, specify the credentials for the user account to be used when using ONTAP snapshots.

For more information about configuring, creating, and using Recovery Points with CDR, see Recovery Points.

BACKING UP DATA FROM ONTAP SNAPSHOTS

Recovery Points created using ONTAP Snapshot can be backed up. For more information about performing backups of Recovery Points, see Backups of Recovery Points.

For information on restoring data that was backed up using ONTAP Snapshot, see Restore Data - ONTAP Snapshot.

RECOVERING DATA FROM ONTAP SNAPSHOT

Data from Recovery Points created using ONTAP Snapshots can be recovered in a variety of ways. For information on recovering data from ONTAP Snapshot, see Recover Replicated Data.

CONSIDERATIONS

Before using CDR with ONTAP Snapshot, review the following information:

GENERAL

Consider the following, if snapshots are created when a LUN is cloned:

- Delete snaps in the reverse order they were created in. If you have a situation where the busy snap is no longer mounted but is still shown as busy, then all additional snaps on this volume created while that snap was mounted will need to be deleted so that this snap will no longer be busy. See the ONTAP 7.3 note below to avoid this dependency.
- Do not mount a volume and create another snap for the volume. To avoid this snapshot dependency, do not manually create a snapshot of a volume while you have a snapshot mounted.

For NetApp ONTAP version 7.3, there is an option to enable the system to only lock backing Snapshot copies for the active LUN clone. If you do this, when you delete the active LUN clone, you can delete the base Snapshot copy without having to first delete all of the more recent backing Snapshot copies.

This behavior in not enabled by default; use the snapshot_clone_dependency volume option to enable it. If this option is disabled, you will still be required to delete all subsequent Snapshot copies before deleting the base Snapshot copy.

We recommend that you enable this option but if you are using any other applications on the LUN, review the documentation for this feature for other impacts. As with this option enabled, if you delete the snapshot that had originally cloned the LUN, then you cannot use "snap restore" to restore the clone from one of the later snaps. If this volume options is later turned off on the volume, then you may have difficulties deleting snapshots because the dependencies will again be enforced.

ISCSI LUN

CDR on Windows supports iSCSI LUNs as destination for NetApp filers. When iSCSI LUN is used as destination to create Recovery Points, you must specify user authentication details. The user information must be specified either in the General tab of the replication set properties on the source or in the General tab of the agent properties on the destination.

The user should have permission to access the filer.

WINDOWS

- One snapshot is created per CIFS share. If `E:` is mapped to `\\me\CDR_1` and `D:` is mapped to `\\me\CDR_2`, there will be two snapshots created on the file server even if both UNC paths end up on the same volume. It may be useful, using this same example, to instead map `D:` to `\\me\CDR_1\D` and `E:` to `\\me\CDR_1\E`, so that only 1 snapshot is needed.
- Snapshots on the NetApp file server will be named `cvcdr_<timestamp>_<number>` , for example `cvcdr_1189118004_1`. If there are multiple shares used, then one snapshot will be created for each share and the `<number>` will be incremented for each. The corresponding CIFS share on the snapshot will be named `1189118004_1`. This share will be in exactly the same place as the share being replicated to, meaning for example that if `\\me\CDR_1` is for path `/vol/myvol/dir1/dir2`, then CIFS share `1189118004_1` will be on `/vol/myvol/.snapshot/ cvcdr_1189118004_1/dir1/dir2`.
- NetApp file server destination should be a NTFS qtree.

UNIX

- It is recommended that you add the local mount path of the filer to the list of FS to be auto-mounted on reboot. For instance, on AIX you need to add it to `/etc/filesystems` file and set the `Mount` to true. This will ensure that if the destination machine gets rebooted, the source machine will continue to replicate to the destination path on the filer and all the CRP snapshots will be automatically mounted.
- To avoid the creation of additional folders on the destination machine, it is recommended that you turn on the nosnapdir option on the volume. To do so, run the `vol options <volname> nosnapdir` command on the ONTAP console.

LICENSE REQUIREMENT

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

Replication Policy

Topics | How To | Related Topics

OVERVIEW

Replication Policies provide a means through which you can configure multiple replication sets within a CommCell from a centralized template. Replication policy consists of a common configuration for a replication set that can be applied to newly created replication set within the CommCell.

During the Replication Policy creation, you configure the template for a replication set by defining the replication type, replication destination, common base folder, and other replication options. Optionally, you can also add replication pairs to the policy. Replication sets (also called Target Replication Sets) created from these replication policies inherit the properties specified in the replication policy. If the replication policy contains any replication pairs, then the replication set is created with the replication pairs by default. Once created, the properties of the replication set can be modified; additional replication pairs can be added and existing pairs can be deleted.

CONFIGURABLE PROPERTIES

This section explains the configurable properties of a Replication Policy.

Once created, the properties of a Replication Policy can be modified at any later time. Any changes to a Replication Policy will affect only the Replication Sets that are newly created after that point.

GENERAL

DESTINATION HOST

Select a common Destination Host for all Replication Pairs within the Replication Set. All replication sets created from this replication policy will have the same Destination Host.

COMMON BASE FOLDER

Common Base Folder, which is the destination paths for all Replication Pairs in the Replication Set. This can be useful when you are adding multiple Replication Pairs at the same time.

For example, if you specify `\replication\` as the base directory, your Replication Pairs might use the following as their destination directories:

```
\replication\rep_pair_1\
\replication\rep_pair_2\
\replication\rep_pair_3\
```

DATA REPLICATION TYPE

Select the replication type for the replication policy. One replication policy can have only one replication type. The following three types of replication is possible:

- Continuous Replication **-** to create replication sets for CDR to replicate data in Continuous Replication mode. See ContinuousDataReplicator for more

information.

- Disk Library Replication **-** to create replication sets for CDR to replicate data in Disk Library Replication mode.
- Workstation Backup **-** to create replication sets for Workstation Backup Agents.

PRE/POST PROCESS

Specify commands for Pre/Post Processes for Data Replication Activities in the Pre/Post Process tab.

REPLICATION OPTIONS

Provide additional options for data replication in the Replication Options tab.

ORPHAN FILES

Specify how Orphan Files will be handled, when discovered on the destination computer. For more information, see Orphan Files. The following default settings apply:

- Continuous Replication - Orphan file handling is turned On by default.
- Workstation Backup - Orphan file handling is turned Off by default.

ADVANCED

Specify the advanced hashing options, enable Disk Library Replication, and set mining tool options (applicable for ContinuousDataReplicator) in the Advanced tab.

FILTERS

Specify files, folders, and patterns to be excluded from data replication activities. For more information, see Replication Filters.

REPLICATION POLICY OPERATIONS

The following operations can be performed on a Replication Policy.

CREATE A REPLICATION POLICY

You can create a new Replication Policy from the **Create New Replication Policy** dialog box. The name of the Replication Policy, Destination Host, Common Base Folder, and the Data Replication Type in the **General** tab are essential components for a Replication Policy. Follow the options in the rest of the tabs and provide the details as applicable. See Create a Replication Policy for step-by-step instructions.

ADD/UPDATE REPLICATION PAIRS

You can add replication pairs to the replication policy, defining the source path from which data is replicated. When defining the source path, keep in mind that the path specified in the replication pair must be a valid path in the client computer in which the replication pairs are created using this replication policy.

See Add/Update Replication Pairs to a Replication Policy for step-by-step instructions. You can add multiple replication pairs to a replication policy. Replication sets created with this policy will contain these replication pairs upon creation.

DELETE A REPLICATION POLICY

When a replication policy is no longer required, you can delete the replication policy. See Delete a Replication Policy for step-by-step instructions.

APPLY SETTINGS

Replication Set details for multiple Workstation Backup Agents can be simultaneously updated using the Apply Settings option. You can create a Replication Policy with the desired settings, or modify the details of an existing Replication Policy, and use the Apply Settings option to apply the Replication Policy configuration to the Replication Sets to selected Workstation Backup clients. When applied, the details of Replication Sets of the corresponding Workstation Backup client computers are modified, the Replication Pairs are reset to New Pair state, and the subsequent backup is converted to a Full Backup.

> **WARNING:**
>
> Using the Apply Settings option applies the Replication Policy settings to the Replication Sets, resets all the corresponding Replication Pairs to **New Pair** state, and converts the following backup job to a **Full Backup**. It is recommended that you use this option with utmost caution.

See Apply the Replication Policy Settings to Replication Sets for step-by-step instructions.

Back to Top

Replication Policy - How To

Topics | How To | Related Topics

CREATE A REPLICATION POLICY

To create a Replication Policy:

1. From the CommCell Browser, right-click the **Replication Policies** node, then click **New Replication Policy**. The Create New Replication Policy dialog is displayed.
2. In the General tab, specify a name for the Replication Policy and a Destination Host that has the replication Agent installed. If required, for Windows, in the **User Name** field, specify a Windows User Account that is part of the Local Administrator group on the destination, with read/write permissions for the product installation folders on the destination. If the agent uses replication logs, and the replication logs directory on the destination is not in the default location, then the **User Name** supplied here must also have read/write permissions for that directory as well.
3. Specify a Common Base Folder so that all destination paths for all Replication Pairs in the Replication Sets (created with this Replication Policy) will be created in the same location on the destination host.
4. Specify the type of replication for which the replication sets created using this replication policy will be used. The following replication types are available.
   - **Continuous Replication -** to create replication sets for CDR to replicate data in ContinuousDataReplication mode.
   - **Disk Library Replication -** to create replication sets for CDR to replicate data in Disk Library Replication mode.
   - **Workstation Backup -** to create replication sets for Workstation Backup Agents.
5. In the Pre/Post Process tab, specify any commands to be run before a data protection operation (like Recovery Point creation, backup job, etc.) and/or after the operation, and specify a User Account with permissions to run the command(s) on the destination machine. For more information, see Pre/Post Processes.
6. In the Replication Options tab, make any appropriate selections:
   - Enable or Disable Software Compression for a Replication Set
   - Data Integrity Validation on transfer (See Enable (or Disable) Data Integrity Validation on Network)
   - Configure the Replication Set for Data Encryption
   - Configure CDR Recovery Points
   - Configure CDR for Backups of Recovery Points
7. In the Orphan Files tab, specify how to treat data on the destination, which is not on the source machine.
8. Optionally, filter the content of any Replication Pairs that are created for this Replication Set in the Filters tab. For more information about Filters, see Replication Filters. To exclude data from data replication operations, do one of the following:
   - Click the **Add** button and, in the **Enter Path** window, type the complete path (including drive letter) of the file/folder/directory that you want to exclude from data replication. Repeat this step if you want to add more files/folders/directories to the filter.
9. Click **OK** to save your changes.

   User Security is configured at the Replication Set level. To specify the security options, expand the Agent icon of the source machine, right-click the Replication Set and select **Properties**, then select the Security tab. For more information, see User Administration and Security.

ADD/UPDATE REPLICATION PAIRS

Required Capability: Capabilities and Permitted Actions

To add new Replication Pairs, or update existing Replication Pairs:

1. In the CommCell Browser, right-click a Replication Policy to which you wish to add the Replication Pair, and select **All Tasks**, then select **Manage Pairs**.
2. To add a Replication Pair, in the Manage Pairs dialog, click **Add**. In the Add Replication Pair dialog box, type a source volume or directory for the Source

Path. Keep in mind that this path must be valid source path in the target computer where the replication pair is created, using this replication policy.

3. To edit a Replication Pair, in the **Manage Pairs** dialog, select the Pair to edit, click **Edit.** Make any changes to the source path in the **Edit Replication Pair** dialog box, and click **OK**.
4. To specify that all VSS writers must be engaged when the acquiring the source snap, select the **Engage All VSS Writers** option. Use this option to create backups that are consistent with the source. The **Automatically detect mount points** option is not applicable when adding Replication Pairs to the Replication Policy.
5. When your changes are complete, click **Apply**.

Filtering of content for Replication Pairs is configured at the Replication Set level; see Replication Filters.

CREATE A REPLICATION SET FROM A REPLICATION POLICY

Before You Begin

- Review Replication Policy.

To create a Replication Set from a Replication Policy:

1. In the CommCell Browser, right-click the Agent icon of the source machine, select **All Tasks**, then select **Add Replication Set From Policy**.
2. In the **Select a Policy** dialog box, select the replication policy based on which you wish to create the replication set. The replication policies created for the corresponding agent will be displayed. Click **OK**.
3. In the **Enter Name** dialog box, specify the name of the replication set, and click **OK**.
4. The replication set will be created with the configurations available in the selected replication policy.

DELETE A REPLICATION POLICY

Required Capability: Capabilities and Permitted Actions

To delete a Replication Policy:

1. From the CommCell Browser, select the **Replication Policies** node, right-click the desired Replication Policy, click **All Tasks** and then click **Delete**.
2. Click **Yes** to the confirmation message.
3. The Replication Policy will be deleted. Note that Replication Sets created from this Replication Policy will have no impact upon deletion.

APPLY THE REPLICATION POLICY SETTINGS TO REPLICATION SETS

Required Capability: Capabilities and Permitted Actions

To apply the Replication Policy settings to Workstation Backup agent Replication Sets:

1. From the CommCell Browser, select the **Replication Policies** node, and select the desired Replication Policy. If you wish you modify the details of the Replication policy, right-click the Replication policy and select **Properties**. Modify the details as required.
2. Right-click the Replication Policy, click **All Tasks** and then select **Apply Settings**. The Apply Settings dialog box is displayed with the list of Workstation Backup client computers configured to the CommCell.
3. Select the Workstation Backup Client computers you wish to apply the Replication Policy settings. Use **Select All** and **Clear All** options to select or clear all the clients in the list.
4. Click **OK** to apply the settings.

Back to Top

Monitoring Data Replication

Topics | How To | Related Topics

OVERVIEW

Replication is a continuous activity and details of on-going replication activity is shown in the Data Replication Monitor in the CommCell Console. See View Data Replication Monitor for step-by-step instructions.

From the Replication Monitor you can:

- View details of data replication activities. See View details of data replication activities for step-by-step instructions.
- View the failed files for a Replication Pair. See View the failed files for a Replication Pair for step-by-step instructions.
- Filter which clients activities are displayed. See Filter which clients activities are displayed for step-by-step instructions.
- Send Log Files of Replication Pair. See Send Log Files for step-by-step instructions.
- Start/Suspend/Resume/Abort data replication activity. See Start/Suspend/Resume/Abort Data Replication Activity for step-by-step instructions.

All other job-based activity, such as Recovery Point creation, is reflected in the Job Controller. See Controlling Jobs in Job Management for comprehensive information.

JOB PHASES

CDR utilizes phases to perform three types of operations - initial data transfer or baselining, smart synchronization, and continuous data replication. The sequence of these phases is listed below along with details of CDR activities during each phase, and the consequence of an interruption, such as a temporary loss of connectivity:

JOB PHASE	ASSOCIATED ACTIVITY	COMMENTS
Baseline Scan	For Windows only, start NTFS journaling on the source to track any file operations that occur during the entire Baseline phase. Scan source path to obtain the number of files and bytes to transfer. Generate Collect File.	The Replication Pair will show a Job State of **Preparing for Replication** in the Data Replication Monitor. If this phase is interrupted: • For Windows, it can resume again at the same point. • For Unix, it will start over. A Full Re-Sync will start at this phase.
Baseline (For Windows)	Calculates checksum on the source and destination to identify files that will be sent to the destination. Data is transferred from the Replication Pair source path to the destination path using the checksum.	If this phase is interrupted, it can resume again at the same point.
SmartSync Scan	Create a non-persistent snapshot; for Windows, compare it to the change journal. Scan snapshot and generate a new Collect File for any files or directories that were *added* or data that was *modified* since the beginning of the Baseline Scan phase.	For Windows: • During the SmartSync Scan phase, CDR requires a short period of no disk activity on the source to begin monitoring the source path, and if there is significant I/O on the source, this can fail, although CDR will continue making successive attempts to find a short period of inactivity. When you have multiple, active Replication Pairs on the same source computer, CDR requires this short period of no disk activity for

		all of them, even for Replication Pairs that use different drives as their source. For example, if you have a Replication Pair on the Client which is already replicating, using `F:\` as its source, and you create a new Replication Pair on the same Client using `G:\` as its source, too much I/O on `F:\` can cause CDR to fail to begin monitoring `G:\`. For Unix: • To create the non-persistent snapshot on the source, CDR requires a short period of time, during which no files are deleted, created, or renamed. (File writes will not affect the snapshot.) If one of these three operations occurs while the snapshot is being taken, the process of creating the snapshot will begin again. If this phase is interrupted: • For Windows, it can resume again at the same point • For Unix, it will start over A Smart Re-Sync will start at this phase.
Processing Orphan Files	Compare the Collect File to the Destination to identify orphan files, and apply orphan file settings.	Any data that was *deleted* on the replication source during the Baselining phases are treated according to your settings for Orphan Files. If this phase is interrupted: • For Windows, it will resume again from the beginning of this phase; however, if the snapshot is no longer available, it will return to the SmartSync Scan phase. • For Unix, it will return to the SmartSync Scan phase.
Checksum Calculation (On Windows only)	Calculate checksums on the source and destination to identify files that have changed since Baseline Scan.	If this phase is interrupted, it will resume again from the beginning of this phase; however, if the snapshot is no longer available, it will return to the SmartSync Scan phase.
SmartSync	Transfer all changed files to destination from the new Collect File.	If this phase is interrupted: • For Windows, it will resume again from the beginning of this phase; however, if the snapshot is no longer available, it will return to the SmartSync Scan phase. • For Unix, it will return to the SmartSync Scan phase.
Updating Smart Sync (On Windows only)	Compare time stamps on source and destination and update. Temporary snapshot is deleted.	If this phase is interrupted, it will resume again from the beginning of this phase.
Replication	Data is continuously replicated from the source to destination.	Log Transfer & Log Replay activity is on-going. For more information, refer to Replication Logs. The Replication Pair will show a Job State of **Replicating** in the Data Replication Monitor. If the Replication phase is interrupted, when restarted, if it is possible, replication will begin again from the last log replayed on the destination; if this is not possible, the Replication Pair will return to the Baseline Scan phase (Full Re-Sync) or to the SmartSync Scan phase (a Smart Re-Sync) depending on the nature and duration of the interruption. Note that if a user manually restarts Replication by choosing **Start Full Resync**, the Replication Pair will return to the Baseline Scan phase.

For the SmartSync Scan, while *new* files and directories will be copied in their entirety, *modified* files do not need to be copied. Thus, for larger files, only the modified portion is re-copied, while smaller files with substantial changes may be copied in their entirety. Modified files below a certain size threshold are copied again as complete files, while files above that size are broken into blocks with just the changed blocks copied to the destination computer.

Files smaller than 256KB will be copied in their entirety whether they match the destination or not. For files above 256KB in size, only the changed blocks will be transferred; the default block size for hashing is 64KB. The default values of the minimum file size and the block size for hashing, can be configured in Replication Set Properties. See Create a Replication Set for step-by-step instruction.

INTERRUPTIONS AND RESTARTS

By default, CDR handles interruptions by seamlessly restarting replication, but if that is not possible, Smart Re-Sync will be started. However, some interruptions will require a Full Re-Sync. The following sections describes each phase and restart behavior when the phase is interrupted:

SMART RE-SYNC

Smart Re-Sync is the default behavior of CDR when activities are interrupted and cannot be seamlessly restarted at the same point again. In general, CDR endeavors to do the following in such cases, wherever possible:

- continue logging on the source
- continue replaying logs on the destination which were received before the interruption
- restart activities exactly where they were interrupted, or as close to that point as possible

For examples of commons types of interruptions, and how Smart Re-Sync handles the recovery, refer to System Behavior when Replication is Interrupted.

For a detailed listing of each phase, and the specifics of the exact point at which Smart Re-Sync restarts activities, refer to Job Phases.

FULL RE-SYNC

Full Re-Sync should be necessary only in cases such as the following:

- the data on the destination is altered by means outside of the replication process, e.g., manually deleted or modified, etc.
- an interruption is of long enough duration that the logs overflow on the source

In such a case, all existing content in the destination path is considered inconsistent and Full Re-Sync is recommended to rebuild it again based on the current data in the specified source path. When you start replication from the Replication Set or Replication Pair level, you can specify Full Re-Sync, causing the Replication Pair to begin at the Baseline Scan phase.

Data Replication will be interrupted if a hard disk used for either a source or destination is put into the 'standby' state through the power schema configuration. It will be necessary to abort activity for all affected Replication Sets and restart them again using **Start Full Resync** after such an event.

CHANGES THAT INTERRUPT DATA REPLICATION

Changes to the following configuration items will not be effective until data replication activity has been interrupted and restarted:

- **Job Results Directory** in Client Computer Properties (Job Configuration) - any Replication Pairs in the Replicating state must be aborted and restarted.
- **Impersonate User** in Client Computer Properties (Job Configuration) on a Destination computer - restart destination computer. (This applies on Windows only.)
- **Automatically delete Orphan Files** in Replication Set Properties (Orphan Files) - any Replication Pairs in the Replicating state must be aborted and restarted.
- **Exclude these Files/Folders/Patterns** for content in Replication Set Properties (Filters) - any Replication Pairs will be aborted and restarted by the system.

The following will require data replication to be interrupted and restarted:

- On Windows, if `chkdsk` is run on a hard disk used for either a source or destination, the affected Replication Pairs in the Replicating state must be aborted and restarted using Smart Re-Sync.
- By default, CDR will always replicate only the new or updated data in the source path. If data is deleted on the destination, since there has been no change on the source, that data will not be replicated again, unless you abort the Replication Pair and perform the following to recopy the data from the source to the destination again:
  - On Windows, perform a Full Re-Sync.
  - On UNIX, perform a Smart Re-Sync.

SYSTEM BEHAVIOR WHEN REPLICATION IS INTERRUPTED

There are several ways in which data replication activity can be interrupted, and CDR recovers from each of them in a similar manner. The table below provides a listing of common causes of interruption, and the effect of them on Baselining, SmartSync, and data replication, as well as how CDR recovers from them.

INTERRUPTION	EFFECT OF INTERRUPTION & SMART RE-SYNC
Abort a Replication Pair during Baselining phases	Baselining activities stop on the source. When the Replication Pair is restarted, Baselining activities will resume, restarting at the beginning of the phase if necessary, then SmartSync and data replication activities will begin automatically.
Abort a Replication Pair during SmartSync phases	Logging stops on the source. When the Replication Pair is restarted, SmartSync activities will resume, restarting at the beginning of a phase if necessary, and data replication activities will begin automatically.
Abort a Replication Pair during Replication phase	Logging stops on the source. When the Replication Pair is restarted, for NTFS or UNIX, Smart Re-Sync will continue the data replication activities automatically; for FAT file systems, Full Re-Sync will be necessary.
Suspend a Replication Set	Baselining, SmartSync, and data replication activities stop for all Replication Pairs, but any logging activities will continue on the source.

	When the Replication Set is resumed: • For any Replication Pairs that were performing data replication, CDR will transfer the accumulated logs to the destination, and data replication will continue. • For Replication Pairs that were in the Baselining or SmartSync phases, how activities begin again will depend on the exact phase the Replication Pairs were in, as well as the operating system type.
Graceful or non-graceful shutdown of the source computer	The destination computer continues to replay the logs it has received. When the source computer and software are running again, Replication Pair(s) will be in the **System Aborted** state for some time, then Smart Re-Sync will be performed.
Graceful or non-graceful shutdown of the destination computer	Logging continues on the source. When the destination computer and software are running again: • For any Replication Pairs that were performing data replication, CDR will transfer the accumulated logs to the destination, and data replication will continue. • For Replication Pairs that were in the Baselining or SmartSync phases, how activities begin again will depend on the exact phase the Replication Pairs were in, as well as the operating system type.
CDR software shutdown on the source	All CDR-related activities stop. When the software is restarted, CDR will start Smart Re-Sync.
CDR software shutdown on the destination	Logging continues on the source. • For any Replication Pairs that were performing data replication, CDR will transfer the accumulated logs to the destination, and data replication will continue. • For Replication Pairs that were in the Baselining or SmartSync phases, how activities begin again will depend on the exact phase the Replication Pairs were in, as well as the operating system type.
Replication Service is stopped on the source	Baselining, SmartSync, and data replication activities stop for all Replication Pairs, but logging continues on the source, and the destination computer continues to replay the logs it had received before the service was stopped. When the Replication Service is started again: • For any Replication Pairs that were performing data replication, CDR will transfer the accumulated logs to the destination, and data replication will continue. • For Replication Pairs that were in the Baselining or SmartSync phases, how activities begin again will depend on the exact phase the Replication Pairs were in, as well as the operating system type.
Replication Service is suspended on the destination	Baselining, SmartSync, and data replication activities stop for all Replication Pairs, and log replay stops on the destination, but logging continues on the source. When the Replication Service is started again: • For any Replication Pairs that were performing data replication, CDR will transfer the accumulated logs to the destination, and data replication will continue. • For Replication Pairs that were in the Baselining or SmartSync phases, how activities begin again will depend on the exact phase the Replication Pairs were in, as well as the operating system type.
Interruption of network connectivity (source and/or destination)	Baselining, SmartSync, and data replication activities stop for all Replication Pairs, but logging continues on the source, and the destination computer continues to replay the logs it had received before the network connectivity was interrupted. When network connectivity is restored: • For any Replication Pairs that were performing data replication, CDR will transfer the accumulated logs to the destination, and data replication will continue. • For Replication Pairs that were in the Baselining or SmartSync phases, how activities begin again will depend on the exact phase the Replication Pairs were in, as well as the operating system type. If the network interruption is for a significant amount of time, the following will occur: • On Windows, the status of the Replication Pair will become **Failed**, and will need to be restarted manually with Smart Re-Sync when connectivity is restored. • On UNIX, CDR will continue to retry sending the logs to the destination computer until network connectivity is restored.
Source computer runs out of log space (Windows) -- or -- Source computer tries to create new entries in a log before the old entries have been transferred to the destination (UNIX)	Logging will stop, all logs will be deleted, all Replication Pairs will be System Aborted. • On Windows, the system will wait 3 minutes, then check space on the log volume. If there is sufficient space, a Smart Re-Sync will occur; if not, the Replication Pair will be Aborted. • On UNIX, a Smart Re-Sync will occur.

For instructions on restarting replication after it has been interrupted, see Start/Suspend/Resume/Abort Data Replication Activity.

JOB STATES

The Data Replication Monitor shows the state of each Replication Pair. These states are briefly described:

New Pair	The Replication Pair has been created, but no activity has taken place yet.
Preparing for Replication	CDR is scanning the source paths, preparing for initial transfer or Full Re-Sync.
Baseline	For detailed information, see Baseline.

Initial Sync	For detailed information, see Baseline Scan.
SmartSync Scan	For detailed information, see SmartSync Scan.
SmartSync	For detailed information, see SmartSync.
Processing	For detailed information, see Processing Orphan Files.
Replicating	Data is being continuously replicated.
Replicating (Not verifiable)	The most recent communication between the CommServe and CDR Client indicated the job was in the Replicating state, but this cannot be verified because communication has been interrupted.
Suspended	Replication activity has been temporarily halted, either by a user, or because communication between the source and destination has been interrupted. Logs continue to be written on the source.
Pending	There has been a temporary interruption and CDR is attempting to reconnect and resume operations.
Failed	Phase failed to complete, or log transfer has stopped, perhaps for connectivity issues; logs continue to be written on the source.
Paused	CDR is trying to resume replication activity.
Stopped	Replication activity has been halted by one of the following: • the Replication Log is inaccessible to be read from on the destination computer • the system, because communication between the source and destination has been interrupted and cannot be successfully resumed by the system • the system, because the replication destination has run out of space • a source or destination throttling condition has not been cleared; if the condition is on a source computer, the job will be System Aborted at first, and if on a destination computer, the job will be Paused at first
System Aborted	For CDR on Windows only, a Replication Pair will be in this state for 3 minutes if the source disk hosting replication logs runs out of space, after which the system will attempt to restart.

To see more information about a particular Replication Pair, see View details of data replication activities.

You can change the state of a Replication Pair, or several at the same time. See Change the State of Replication Pair.

CONSIDERATIONS

- The status of all Replication Pairs is not immediately updated when one Replication Pair is resumed. For instance, when all Replication Pairs had been placed in the Paused state, and you Resume one of them, a prompt will ask if you want all Pairs to be resumed. If you choose to do so, all the Replication Pairs that were placed in the Paused state will Resume, and be placed back in the same state they were in previously. However, the CommCell Console will not immediately reflect the status of all the other Replication Pairs that were resumed, and they may still be shown in the Paused state. The CommCell Console will properly synchronize and display the correct state of the Replication Pairs within a few minutes.
- During SmartSync of application data, Data Replication Monitor may display more than the actual number of files transferred.

JOB DETAILS

The following information is available in the Data Replication Monitor:

Active	When the symbol is green, it indicates recent activity for the Replication Pair; an orange symbol indicates no recent activity. An exclamation point preceding the symbol indicates that some files are not copied successfully to the destination computer during replication. To see failed files for a replication pair, see View the failed files for a Replication Pair for step-by-step instructions.
Phase	The current phase of the job; for more detailed information see Job Phases.

GENERAL

Job ID	A unique number allocated by the Job Manager for the operation.
State	The current state of the Replication Pair; for more detailed information see Job States.
Last Update Time	The date and time of the CommServe when the Job Manager last updated the Data Replication Monitor.
Pair Abort Reason	For a Replication Pair that was aborted, the reason is listed.
Last Error	The most recent error message for this Replication Pair.

INITIAL SYNC INFORMATION

Start Time	The date and time of the CommServe when data replication activity began for the Replication Pair.
Number of Files To Be Transferred	The files remaining to be transferred for the Replication Log file currently being replayed on the destination.
Number of Files Already Transferred	The files transferred for the Replication Log file currently being replayed on the destination.
Data To Be Transferred during Initial Sync On Source	The aggregate size of all files to be transferred between the source and destination for the Replication Pair. The actual data transferred may differ slightly from this number, based on whether a given file actually gets transferred in full or in part.
Data Transferred during Initial Sync On Destination	The sum of all data already transferred between the source and destination for the Replication Pair.
Throughput Unit	The rate of data transfer during Baseline phase, in GB/hour.
Progress	The percentage of files transferred for the Replication Log file currently being replayed on the destination.

REPLICATING STATE INFORMATION

Last Log Played Time	The date and time of the CommServe when the most recent Replication Log was played on the destination computer.
Replicated Data	The sum of all data transferred between the source and destination machines since the Start Time.
Attempts	The number of attempts at replication the system has made for the Replication Pair.
Latest Source Log	The number of the most recent Replication Log that was created on the source computer.
Latest Destination Log	The number of the most recent Replication Log that was replayed on the destination computer. If this number is lower than the Latest Source Log number, it indicates that the destination computer has not yet replayed all of the Replication Logs that have been created on the source computer.

CONFIGURATION

Pair ID	A unique number allocated by the Job Manager that identifies the Replication Pair.
Source Path	The path on the source computer for the Replication Pair.
Destination Path	The path on the destination computer for the Replication Pair.
Replication Set	The name of the Replication Set.
Replication Type	The type of replication configured for the Replication Set. (See Data Replication Type.)
Client	The CDR Client that is the source computer for the Replication Pair.
Destination Host	The CDR Client that is the destination computer for the Replication Pair.

ATTEMPTS

The following information is available in the Attempts window:

Phase	The phase that the Replication Pair was in at the time of the attempted activity.
State	Current state of the Replication Pair.
Start Time	The date and time of the CommServe when the attempted activity began for the Replication Pair.
End Time	The date and time of the CommServe when the attempted activity ended for the Replication Pair.
Elapsed Time	The amount of time that elapsed while the activity was being attempted for the Replication Pair.
Files to Transfer	Files to be transferred to the destination computer for the Replication Pair, based on the initial scan.
Files Transferred	Files already transferred to the destination computer for the Replication Pair.
Data Transferred	The sum of all data already transferred between the source and destination during the attempted activity.
Data to Transfer	The aggregate size of all files to be transferred between the source and destination for the Replication Pair. The actual data transferred may differ slightly from this number, based on whether a given file actually gets transferred in full or in part.

Back to Top

Monitoring Data Replication - How To

Topics | How To | Related Topics

View Data Replication Monitor

View details of data replication activities

View the failed files for a Replication Pair

Filter which clients activities are displayed

Send Log Files

Start/Suspend/Resume/Abort Data Replication Activity for a Replication Pair

Suspend Data Replication Activity from the Replication Set

Resume Data Replication Activity from the Replication Set

VIEW DATA REPLICATION MONITOR

Before you begin

- Review Monitor Data Replication

Required Capability: Capabilities and Permitted Actions

To view data replication monitor:

1. In the CommCell Browser, right-click the CDR icon, and select **All Tasks**, then select **Data Replication Monitor**.

In the Data Replication Monitor, all replication activities will be displayed for each Replication Pair. For more information about the States displayed, see Job States.

VIEW DETAILS OF DATA REPLICATION ACTIVITIES

Required Capability: Capabilities and Permitted Actions

To view details of data replication activities for a Replication Pair:

1. In the Data Replication Monitor, right-click any Replication Pair, and select **Details**.
2. The Pair Activity will be displayed with additional details about the job.

VIEW THE FAILED FILES FOR REPLICATION PAIR

Required Capability: Capabilities and Permitted Actions

To view the failed files for a Replication Pair:

1. In the Data Replication Monitor, right-click any Replication Pair, and select **View Failed Files**.
2. If any failed files have been logged for this Replication Pair, they will be displayed.

FILTER WHICH CLIENTS ACTIVITIES ARE DISPLAYED

Required Capability: Capabilities and Permitted Actions

To filter which clients' activities are displayed:

1. In the Data Replication Monitor, right-click any Replication Pair, and select **Filters**.
2. In the Filter Operations dialog box, select from the following:
   - Select an existing filter from the **Filters** list. (You can click the **View Filter Details** button to see which clients each filter is configured to monitor.)
   - Click the **Create a New Filter** button. In the **Create a New Filter** dialog box, select the clients to be monitored and click **OK**.
3. Click **Apply** and then click **Close**.
4. The Data Replication Monitor will now show activities only for the Replication Pairs for the clients defined in your selected filter.

SEND LOG FILES

Required Capability: See Capabilities and Permitted Actions

To send log files:

1. In the Data Replication Monitor, right-click any Replication Pair, and select **Send Logs**. The **Send Log Files** window is displayed.
2. From the **General** tab, provide the following information:
   - one or more recipient email addresses (multiple email addresses separated by a comma)
   - a **Subject (Ticket Number)**
   - a **Problem Description**
3. From the **Computers** tab, select the following options, if desired:
   - **Job ID** if you are sending the log files of a specific job.
   - **App ID** if you are sending the log files of a specific replication pair.
   - If you wish to send the log files to specific computer(s), you can select and move the appropriate client computers from the **Available Computers** list to the **Selected Computers** list.
4. From the **CommCell Information** tab, select the following options, if desired:
   - **CommServe** if you wish to include the CommServe database log files.
   - **SRM Database** if you wish to include the SRM database log files.
   - **Database logs** if you wish to include the error logs for the CommServe Database Engine.
5. From the **Time Range** tab, select the **Time** check box if you wish to include a specific number of hours or days to be included in the log files prior to the current time or date.

Then, select the **Select Date/Time** option if you wish to choose a specific range of dates to be included in the log files. You may also select the time zone that corresponds to the dates and times selected.

6. From the **Machine Information** tab, select the log information you wish to include, if any.
7. From the **Output** tab, select the following options:
   - If you wish to upload the log files to an existing FTP and/or HTTP location, select the **Upload to FTP Location** and/or **Upload to HTTP Location** check box.
   - Click the **Proxy Settings** button. This opens the Proxy Setting dialog box. Specify Proxy Location and Proxy Port number.
   - If you wish to send the log file information immediately, select the **Email to Recipients** box. Type the address of the recipient(s) who are to receive the log file information.

     To send log files to multiple recipients, select the desired recipients in the **Available Users** field. You can then add or remove the desired recipients from the **Users to be Notified** using the appropriate arrow buttons. Note that log files can be configured to be sent via e-mail to user groups created from within the CommCell Console as well as external domain user groups.
   - To save the log files to a specific directory on the CommServe disk, select the **Save to Folder** check box. You can either type or browse the CommServe local path.
8. Click **Ok**.

START/SUSPEND/RESUME/ABORT DATA REPLICATION ACTIVITY FOR A REPLICATION PAIR

Before You Begin

- Review Monitoring Data Replication

Required Capability: Capabilities and Permitted Actions

To Start/Suspend/Resume/Abort a Replication Pair:

- In the CommCell Browser, right-click a Replication Pair, and select **Start, Start Full Resync, Suspend**, **Resume**, or **Abort**.

From the Replication Pair level, you cannot Suspend a Replication Pair that is in the Replicating state; instead, do this from the Replication Set level. See, Suspend Data Replication Activity from the Replication Set for step-by-step instructions.

SUSPEND DATA REPLICATION ACTIVITY FROM THE REPLICATION SET

Before You Begin

- Review Monitoring Data Replication

Required Capability: Capabilities and Permitted Actions

To suspend all data replication activity from the Replication Set level:

1. In the CommCell Browser, right-click a **Replication Set** in the source machine, and select **All Tasks**, then select **Suspend Data Transfer**.
2. For Windows, if you **Suspend Data Transfer** for one Replication set, you will be prompted with a list of Replication Set that have the same destination.

   If you click **Yes** all Replication Pairs on that source computer which use the same destination computer will be placed in the Paused state, even those that are in different Replication Sets.

   For UNIX, data replication activities are suspended for all Replication Pairs in the Replication Set.

RESUME DATA REPLICATION ACTIVITY FROM THE REPLICATION SET

Before You Begin

- Review Monitoring Data Replication

Required Capability: Capabilities and Permitted Actions

To resume all data replication activity from the Replication Set level:

1. In the CommCell Browser, right-click a **Replication Set** in the source machine, and select **All Tasks**, then select **Resume Data Transfer**.
2. For Windows, you will be prompted with a list of Replication Set that have the same destination.

   Click **Yes** to resume all Replication Pairs between this source and destination - including the replication pairs in other replication sets.

   Click **No** to resume only the Replication Pairs in this replication set.

For UNIX, data replication activities are resumed for all Replication Pairs in the Replication Set.

Back To Top

Recovery Point Creation History

Topics | How To | Related Topics

You can view the Recovery Point creation history of ContinuousDataReplicator. The **Job History Filter** dialog box allows you to view detailed, historical information about jobs.

Once you have chosen your filter options, jobs that meet the criteria you selected are displayed in the **Job History** window. From this window you can right-click a job and view more detailed information such as the:

- Items that failed during the job.
- Items that were killed.
- Details and events of the job.
- A list of the Recovery Points that were created.

Back to Top

Recovery Point Creation History - How To

Topics | How To | Related Topics

View Job History Details

View the Items That Failed For a Data Protection Operation

View the Events of a Job History

View the Log Files of a Job History

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 ITEMS THAT FAILED FOR A DATA PROTECTION OPERATION

A listing of files and folders that failed is not available for the Quick Recovery Agent, nor the Image Level and Image Level ProxyHost *i*DataAgents. These agents do not perform a file level backup/copy.

To view the list of items that failed for a data protection operation:

1. From the CommCell Browser, right-click the entity whose history of data protection operations you want to view, click **View**, and then click to view a job history.
2. From the Job History Filter dialog box, select the filter options, if any, that you want to apply, and then click **OK**.
3. From the Job History window, right-click the operation whose list of failed items you want to view, and then select **View Failed Items**. The **Unsuccessful Backup Files** window (for DataArchiver Agents, **Items On Which Archive Failed**) displays those items that failed. If no items failed, a message to that effect is displayed.

4. Click **Close**.

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 where 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

Recovery Point Copyback History

Topics | How To | Related Topics

You can view the Recovery Point copyback history of ContinuousDataReplicator. The **Job History Filter** dialog box allows you view detailed, historical information about job.

Once you have chosen your filter options, jobs that meet the criteria you selected are displayed in the **Job History** window. From this window you can right-click a job and view more detailed information such as the:

- Details of the Copyback job
- Events of the Copyback job
- Log files of the Copyback job
- Source Volume
- Destination Volume


The source and destination fields are not supported for Copyback history.

Back to Top

Recovery Point Copyback History - How To

Topics | How To | Related Topics

View Job History Details

View the Items That Failed For a Data Protection Operation

View the Events of a Job History

View the Log Files of a Job History

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 ITEMS THAT FAILED FOR A DATA PROTECTION OPERATION


A listing of files and folders that failed is not available for the Quick Recovery Agent, nor the Image Level and Image Level ProxyHost *i*DataAgents. These agents do not perform a file level backup/copy.

To view the list of items that failed for a data protection operation:

1. From the CommCell Browser, right-click the entity whose history of data protection operations you want to view, click **View**, and then click to view a job history.

2. From the Job History Filter dialog box, select the filter options, if any, that you want to apply, and then click **OK**.
3. From the Job History window, right-click the operation whose list of failed items you want to view, and then select **View Failed Items**. The **Unsuccessful Backup Files** window (for DataArchiver Agents, **Items On Which Archive Failed**) displays those items that failed. If no items failed, a message to that effect is displayed.
4. Click **Close**.

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 where 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

Disk Library Replication

Topics | How To | Related Topics

OVERVIEW

The Disk Library Replication solution is designed to seamlessly replicate data from remote offices to a centralized data center, where it is stored on disk media for ready access. Consolidating all backup data at the centralized data center provides redundancy for disaster recovery, as well as an alternate source for normal data recovery operations. This innovative solution is easily deployed, without any disruption of current data protection operations, and with no impact on your end users at remote offices; no additional software or configuration is required on client computers, as only MediaAgents are involved. Even more importantly, the data from every *i*DataAgent deployed at remote offices can be replicated to the centralized data center, no matter the source of the data -- all supported file systems, databases, and data types.

Significant benefits result when this solution is combined with several other available features; Deduplication performs deduplication at the remote office before data is replicated to the centralized data center, greatly reducing the network load; for information about MediaAgents and data types supported by deduplication, see Deduplication - Support. Replication Compression provides a similar benefit, compressing data at the remote office, replicating it across the network to the centralized data center, then uncompressing it on the destination MediaAgent; finally, Replication Encryption can be employed to ensure the security of data transmitted across non-secure networks, by encrypting the data before transmission and automatically decrypting it on the destination MediaAgent at the centralized data center.

A typical Disk Library Replication solution is shown in the illustration.

Software is deployed on a MediaAgent located at a remote office, and a MediaAgent located at the centralized data center. Thus, there is no impact on end users and client computers, since all operations involve MediaAgents only. Deployment and configuration occur once, and the solution continues to function without any further operator intervention. Once the initial transfer of data has been completed at the time of deployment, all data replication after that point will involve only data that has changed on the MediaAgent, similar in concept to Incremental Backups. This is different from an Auxiliary Copy operation, which is similar in concept to Full Backups, copying all selected data every time it runs, and also expanding all deduplicated data each time.

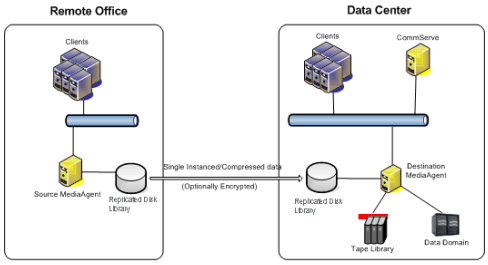

PRE-REQUISITES

The Disk Library Replication solution replicates data from one MediaAgent to another; typically, this would involve a MediaAgent installed in a remote office, and one installed at a centralized data center, although several remote office MediaAgents can replicate their data to the same centralized MediaAgent. Note that each remote MediaAgent will require a different mount path on the centralized MediaAgent.

Each MediaAgent computer requires the following:

- Meets System Requirements for both MediaAgent and ContinuousDataReplicator software. Disk Library Replication solution is also supported on Windows clustered computers.

  The platform on which the software is installed must be supported for both the MediaAgent and ContinuousDataReplicator software components.

- MediaAgent and ContinuousDataReplicator (CDR) software installed; see Installation.
- Shared Disk Library configured to use replicated disks on both source (remote office) and destination (centralized data center) MediaAgents; see Overview of Disk Libraries on Replicated Disks.

In addition, it is recommended that each MediaAgent have the following configured for this solution:

- Deduplication
- Replication Compression

- Replication Encryption (optional, when transmitting data across non-secure networks)

For a list of required product and feature licenses, see License Requirements.

How to Use the Disk Library Replication Solution

This section details how to deploy and configure, temporarily disable, or uninstall the Disk Library Replication solution.

INSTALL, CONFIGURE, AND BEGIN USING THE DISK LIBRARY REPLICATION SOLUTION

1. Install the latest version of the MediaAgent and ContinuousDataReplicator (CDR) software on a source (remote office) computer and destination (centralized data center) computer designated for this solution. For step-by-step instructions, see Installation.
2. Configure the Disk Libraries on Replicated Disks:
   - Using designated disks on the source and destination MediaAgents, Configure Disk Libraries on Replicated Disks; for general information, see Overview of Disk Libraries on Replicated Disks. In the Add Disk Library screen, select the **Enable Replication** option for the Shared Disk Device you just configured. Do not select the **Automatically create storage policy for new datapath** option if you plan to use Deduplication.

     When you have completed configuring the Disk Libraries on Replicated Disks, a Replication Set with the name "SimpleDataReplication" and associated Replication Pair(s) will be automatically created, and will appear in the CommCell Console under ContinuousDataReplicator on the source (remote office) MediaAgent. The Data Replication Type for the Replication Set will be set to Disk Library Replication option. For an example of how Replication Sets and Pairs appear in the Console, see Tree Levels in ContinuousDataReplicator. It is strongly recommended that you do not attempt to change any settings in the Properties of these automatically created Replication Sets and Pairs, except for those discussed in this section.

   - If you are using Disk Library Replication Solution on clusters, ensure that the Disk Library Replication option is selected in the Replication Set (Advanced tab) on all cluster nodes.
3. Create a Storage Policy to be used for this solution, using the Disk Library on Replicated Disks that you just configured; for general information, see Storage Policies. It is recommended that you also enable Deduplication of Data; for instructions, see Deduplication.
4. Configure Subclients (new or existing) with the content to be backed up and replicated to the centralized data center. Make sure to use the Storage Policy you created in the previous step, which will replicate the data to the Disk Library on Replicated Disks. For instructions, see Create a New Subclient; for more information, see Subclients.
5. Synchronization of data between the source and destination MediaAgents can be configured by scheduling the replication jobs at desired intervals using the job scheduler.

   Use the following steps to schedule the replication job:

   - From the CommCell Browser, navigate to **Client Computers** | ***<source_client_computer>***
   - Right-click the **Disk Library Replication**, point to **All Tasks** and then click **Backup**.
   - Select the desired **Backup Type**.
   - From the **Job Initiation**, click the **Schedule** option and then click **Configure** button.
   - Schedule the job as required, from the Schedule Details dialog box and then click **OK**. See Scheduling for more information.
   - Click **OK** to close the Backup Options dialog box.
6. Additional recommended options for this solution:
   - Enable Software Compression for a Replication Set; for more information see Replication Compression.
   - Configure the Replication Set for Data Encryption; for more information see Replication Encryption.
   - Configure Throttling for CDR Replication Activities; for more information see Data Replication - Throttling.
7. Once the installation and configuration is complete on both the source and destination MediaAgents, replication will begin within the update interval specified earlier in this procedure.

TEMPORARILY DISABLE AND RE-ENABLE THE DISK LIBRARY REPLICATION SOLUTION:

Disk Library Replication solution can be temporarily disabled/re-enabled using one of the following methods:

- Using the `nSuspendSDR` registry key, in the source machine.
  - To disable, create the `nSuspendSDR` registry key and set its Value to "1".
  - To re-enable, set the value of the `nSuspendSDR` registry key to "0" or remove the key.
- Disable any job schedules created for the Disk Library Replication solution.

UNINSTALL/DECONFIGURE THE DISK LIBRARY REPLICATION SOLUTION:

1. Stop (abort) replication activities; create the `nSuspendSDR` registry key (if it does not already exist) and set its Value to "1".
2. Delete the Storage Policy that was created for this solution. For instructions, see Delete a Standard Storage Policy.
3. Delete the Replication Pairs and Replication Sets that were automatically created when the Disk Libraries on Replicated Disks were created.

4. Delete any job schedules created for the Disk Library Replication solution.
5. Deconfigure the Disk Libraries on Replicated Disks that were configured on the source (remote office) and on the destination (centralized data center) MediaAgents. For instructions, see Deconfigure Libraries; for overview information, see Deconfiguring Libraries and Drives.

   To leave the CDR software installed and use it to perform data replication, deselect the Disk Library Replication option in the Replication Set (Advanced tab) and select Continuous Replication for data replication type in the Replication Set (General tab). Cycle the Replication Service (CVRepSvc) on both the source and destination computers, and skip the next step. For more information about Services, see Services.

6. Uninstall ContinuousDataReplicator (CDR) software from both the source (remote office) and destination (centralized data center) MediaAgent computers. For more information, see Uninstalling ContinuousDataReplicator.

How to Recover Disk Library Replication Data

Use the following method to recover data that was backed up from a remote office to a centralized data center.

Browse Data from the source client in the usual manner, but select the centralized MediaAgent from the **Use MediaAgent** list in the Browse Options window. Then Restore Backup Data to the original client at the remote office, or to a location of your choosing. This method might be best suited to cases where data is being recovered to a single computer, as opposed to many computers.

Best Practices

Review the following recommendations when deploying the Disk Library Replication solution:

- It is strongly recommended that you employ Deduplication and Replication Compression as part of this solution, as the benefits in reducing the size of data and speeding throughput across the network are significant.
- When this solution will involve data being transmitted across non-secure networks Replication Encryption is recommended to ensure data security.
- When Replication Compression is enabled, do not also enable Data Compression on the MediaAgent.
- When Replication Encryption is enabled, do not also enable Data Encryption on the MediaAgent.
- While it is possible to change the read-write access for mount paths associated with replicated disks, this should never be done for the Disk Libraries on Replicated Disks used for this solution.

Considerations and Notes

Review the following to gain a better understanding of how this solution operates:

- Replication Pairs will display a state of "Stopped" in the CommCell Console when not actively moving data. This is because the Pair will activate at the interval you specified in the **Media Management Configuration (Service Configuration)** Control Panel (or registry key `SDRPairStartIntervalMins`), replicate data transferring any new or changed data, then enter the "Stopped" state until the interval completes again.
- The software automatically creates filters to exclude the area of disk where the deduplication database is active on the source MediaAgent, so that data from this area is never replicated to the destination MediaAgent. In addition, filters are created to exclude temporary files from replication (`/**/*.tmp` in UNIX and `*.tmp` in Windows.)
- The option to automatically delete orphan files is specified when the Replication Set is created by the software; for more information on this subject, see Orphan Files.
- Locked files will not be replicated; once write activity has ended for a particular file, and it is no longer locked, it will be replicated to the destination at the next update interval. Any data in the process of being backed up to the source MediaAgent in the remote office will not be replicated to the centralized MediaAgent. This means that replicated data can only be restored successfully if the backup that they were a part of had completed before the most recent replication started. Any backups that completed before the start time of the latest replication as shown in the Data Replication Monitor will be restorable.
- Most of the features of ContinuousDataReplicator are not functional when it is deployed for this solution.
- The data from a source (remote office) MediaAgent cannot be replicated to multiple destinations (commonly referred to as "fan-out") using this solution.
- If you are using Auxiliary Copy to read data from the destination, then deferred copy settings must be enabled for the Auxiliary Copy, for the number of days required for the replication to be complete. See Auxiliary Copy with Deferred Copies for details.

License Requirements

The following products and features mentioned in this solution require an available Product License or Feature License in the CommServe:

- MediaAgent
- ContinuousDataReplicator (CDR)
- Deduplication

- Disk Libraries
- Data Encryption

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

ContinuousDataReplicator Disaster Recovery Solution for Building a Standby Exchange Server

OVERVIEW

This document describes the procedure necessary to enable rapid recovery of *Production Exchange Server* on the *Standby Server* using ContinuousDataReplicator (CDR). This procedure involves pre-configuring Exchange Server on the *Standby Server* to eliminate these steps during an actual disaster recovery, reducing the recovery time.

This procedure uses the *Standby Server* as the target. However, an intermediate server could also be used as the CDR target. The diagram depicts the configuration used in this approach. In addition to using the live replicated data, this procedure also uses Consistent Recovery Points (CRPs) to allow for alternative recovery points in the event the live (replicated) data can not be used by Exchange Server due to consistency issues.

This procedure has been validated using Exchange 2003 on Windows 2003 with SP1. Newer versions of Exchange should behave the same in these scenarios.

ADVANTAGES

ContinuousDataReplicator provides a mechanism for asynchronously replicating file system and application data to a remote site. Data is replicated at the byte and file level which provides for a very effective use of available network bandwidth, thus making it ideal for disaster recovery scenarios.

CONFIGURATION

The configuration described in this document involves setting up CDR to replicate to the Standby Exchange server directly. The Standby Exchange machine will replace the production machine by assuming its name but keep its existing IP address. The original DNS entry for the *Standby Exchange Server* will remain but the IP address for the *Production Exchange Server* will be set to the IP address of the *Standby Server*.

The following sections discuss preparing the *Production* and *Standby Servers* as well as setting up replication.

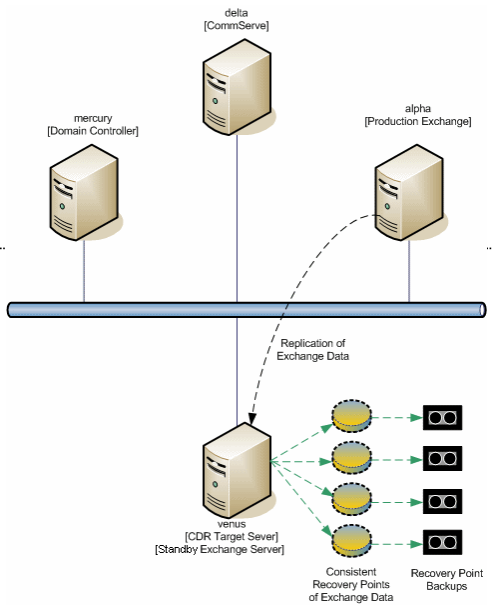

PREPARE THE PRODUCTION SERVER

It is assumed that the *Production Exchange Server* (alpha) is operational and up to date with respect to service packs and/or required patches. Also, it is assumed that the *Standby Exchange Server* (venus), which is also the CDR target, is running as a member of the domain and that the CommServe has already been installed.

The following steps should be followed to prepare the *Production Exchange Server* for CDR:

1. Make sure that the storage groups are moved off of the C: drive. This is usually the case in production environments but the location should be verified.
2. Place the system path on the same volume as the transaction logs. This will make the configuration of the replication pairs simpler since the `.chk` file and the logs will be collocated.
3. Install ContinuousDataReplicator on the server along with the Windows File System ¡DataAgent and the Exchange DB agent. It is expected that standard backups will be executed on the *Production Exchange Server* to manage the logs. Note that CDR will not truncate Exchange logs since all backup operations associated with CDR are file system based. This means that although a CRP will quiesce Exchange and provide for a consistent point in time view, there is no knowledge of the application during the backup phase. Hence, the Exchange APIs which truncate logs are not utilized. Note that as the Exchange DB backup causes logs to be truncated on the production server the deletion of the log files will be replicated to the target by CDR, so cleanup will be automatic on the target server.

   If using the Exchange DB agent is not possible on the *Production Exchange Server* then circular logging can be enabled to manage log growth.

PREPARE THE STANDBY SERVER

Configuring the *Standby Exchange Server* will require that the *Production Exchange Server* be offline while working on the *Standby Server*. These steps could be performed during the recovery time, however, that would increase the recovery time and potentially introduce more opportunity for error. It is recommended that these steps be performed in advance to provide the most straight forward recovery procedure. If the Standby server is being created at recovery time then skip all steps which are performed on the production server.

ON THE STANDBY EXCHANGE SERVER (VENUS)

1. Install ContinuousDataReplicator software.
2. Create volumes and folders for storing Exchange data and logs equivalent to those on the *Production Exchange Server*. These will be used for replication targets for the live data as well as operational data and logs locations when the *Standby Exchange Server* is being used.

ON THE PRODUCTION SERVER (ALPHA)

3. 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**
4. 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
5. Stop All running Exchange Services.
6. Rename the *Production Exchange Server* to a temporary name (oldalpha).
7. Reboot the *Production Exchange Server* so that the new name takes effect.

   At this point Exchange will be off-line and the *Standby Server* can be configured.

ON THE STANDBY EXCHANGE SERVER (VENUS):

8. Rename the *Standby Server* (venus) to original *Production Exchange Server* (alpha).
9. Reboot the *Standby Server.* **The *Standby Server* is now alpha**.
8. Install the following Windows components, if they are not already installed:
   - NNTP (for Windows 2000 & 2003)
   - SMTP (for Windows 2000 & 2003)
   - ASP .NET (for 2003)
9. Create Partitions if necessary and folders to match the locations of the Exchange logs and databases on the production Exchange server.
10. 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`

    - 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`
    - When the installation wizard starts make sure that you install all of the components that were on the *Production Exchange Server*. Note that by default the Microsoft Exchange System Management Tools may not be selected automatically. You will need the System Manager in subsequent steps so make sure that the System Management Tools are selected.
11. Apply the Service Pack(s) and patches that were applied to the production Exchange server.

    Example: `Z:> update.exe /DisasterRecovery`

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.

12. Make sure that all of the Exchange services are stopped.
13. Set all the Exchange services to manual startup.

ON THE DOMAIN CONTROLLER

14. 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:
        - Configuration Container
        - CN=Configuration
        - CN=Services
        - CN=Microsoft Exchange.
    b. Right-click your organization [Domain Name] then click **Properties**. Click the **Security** tab:
        - 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**.
        - 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:
        - CN=[Domain Name]
        - CN=Administrative Groups
        - CN=First Administrator Group
    d. Right-click **CN=Servers** and select **Properties**. Click the **Security** tab:
        - 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**.
        - 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:
        - 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**.
        - Highlight the server object and click **Full Control** under Allow, then click **OK**.

ON THE STANDBY EXCHANGE SERVER (ALPHA, PREVIOUSLY VENUS**)**

15. Manually start the Exchange services on the *Standby Server*.
16. Verify that the Transaction Logs and System Path are located as they were on the *Production Server* using the Exchange System Manager.
17. Also verify that the Stores and Public Folders are correctly located.
18. Stop all of the Exchange services on the Standby server. Make sure to set all startup types to Manual.
19. Rename the *Standby Server* back to its original name (venus) and reboot for the change to take effect. Check the Active Directory to make sure that it has been updated and that the Standby server is listed by its original name.

ON THE PRODUCTION EXCHANGE SERVER (OLDALPHA)

21. The Exchange services on the *Production Exchange Server* can now have their Startup Types reset back to their original values.
22. At this point the *Production Exchange Server* (oldalpha) can be renamed back to its original name (alpha) and rebooted.
23. Using the Exchange System Manager mount the Mailbox and Public Folder stores on the *Production Exchange Server* (alpha).
24. Verify that Exchange is operational and that email / public folders are accessible.

SETTING UP REPLICATION

Now that the servers are prepared and the CommCell software has been installed, the replication set and associated replication pairs can be set up and replication initiated. The replication set and pairs can be either configured manually, or by the use of wizard. Though not required, when live replica data is used to bring up the *Standby* Exchange server - it avoids a Copyback operation. Note the destination paths on the target must be set to mimic the source. See ContinuousDataReplicator for more details on installation and configuration of CDR.

When the pairs are in sync and all of the initial data has been replicated to the target (venus), the next step is to schedule Consistent Recovery Points along with backups of those Recovery Points. This allows for Point-In-Time recovery for the Exchange Server with known consistent data. Alternatively, the live data can be utilized to provide the closest to up to the minute recovery. However, note that live data will be crash consistent at best and is not guaranteed to be recoverable by Exchange.

BRING THE STANDBY SERVER ONLINE

The strategy for replacing the original *Production Server* with the *Standby Server* is illustrated below:

Because of the Exchange dependencies with Active Directory, changes are required for Exchange to operate correctly on the *Standby Server.*

1. It is likely that the initiation of the fail-over to the *Standby Exchange server* is due to the *Production Exchange Server* crashing but in any case the production server (alpha) should be shutdown.
2. On the Domain Controller RESET the *Production Exchange server* computer entry. This forces the computer to rejoin the domain.
3. Remove the *Standby Exchange server* (venus) from the domain and reboot. Removal from the domain and the associated name change appear to have to be done in two steps, which necessitates multiple reboots to get the complete switch to occur.
4. On the Domain Controller using Active Directory Users and Computers delete the entry for the Standby Exchange server (venus).
5. Rename the Standby Exchange server (venus) to the production Exchange server (alpha) and reboot.
6. Rejoin the Standby Exchange server to the Domain as the production Exchange server (alpha) and reboot.
7. Verify/Change the IP address of the *Production Exchange server* (alpha) to match the IP address of the *Standby Exchange server* (venus). By changing the IP address of the original *Production Exchange server* (alpha) to the *Standby server's* address Outlook clients will be able to connect to Exchange correctly. Leave the existing entry for the *Standby Exchange server* (venus) intact so that the system can be accessed by its original name as well.
8. At this point the live replica data can be used or a recovery point selected to recover Exchange. Note that live data is crash consistent at best and can pre-date the crash point in time due to the asynchronous nature of CDR. However it will provide the closest to current data for recovery purposes and not require a copy operation since it is in the proper location already. If recovery from the latest replicated data is not desirable or possible due to inconsistencies then a CRP or backup can be chosen instead. From the CommCell Console browse the CRPs and select the one to be utilized on the *Standby server*.
9. Recovery of CRP data can be done two ways. (Use one method.)
   a. The Copyback button can be used to initiate the copy operation from the GUI. The original *Standby Exchange server* name (venus) is selected as the recovery host and the Copyback function is initiated. This is the most straight forward method.

      NOTE: The copyback will overwrite the live copy so be sure that the live copy will not be required in the future. Otherwise manually copy the data to a new location before issuing the Copyback function. After selecting the *Standby Exchange Server* by its original name (venus) click OK to start the copy operation.

   b. Alternatively the CRP could be mounted and manually copied to the required location. This is accomplished through the View Snapshots function. Click **View Snapshots** Each of the snapshots in the CRP will be listed and each should be mounted by selecting the individual snapshot and clicking the Mount button which exposes the mount point dialog so that a drive letter can be assigned.
10. Start the Exchange services on the *Standby Exchange server.*
11. Start the Exchange System Manager on the *Standby Exchange server* and mount all mailbox and public stores
12. Verify that Exchange is operational and mailbox and public stores are accessible.
13. Optionally, set the following Exchange services to *Automatic* and start these services. (This is advisable if the *Standby Exchange server* is expected to be operational for an extended period of time.)
    - 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

    In addition the **Do not mount this store at start-up** option on the Mailbox and Public Store could also be un-checked so that the stores will mount automatically.

Back to Top

ContinuousDataReplicator Disaster Recovery Solution for Building a Standby SQL Server

OVERVIEW

This document describes the procedure necessary to enable rapid recovery of Standby SQL Server at the disaster recovery site using ContinuousDataReplicator (CDR). In addition to using the live replicated data, this procedure also uses Consistent Recovery Points (CRPs) to allow for alternative recovery points in the event the live (replicated) data can not be used by SQL Server due to consistency issues. This procedure involves pre-configuring SQL Server on the Standby machine to eliminate these steps during an actual disaster recovery, reducing the recovery time.

The configuration described below uses the Standby SQL Server machine as the target. However, an intermediate server could also be used as the CDR target.

This procedure has been validated using SQL Server 2000 on Windows 2003 with SP1. Newer versions of SQL should behave the same in these scenarios.

ADVANTAGES

ContinuousDataReplicator provides a mechanism for asynchronously replicating file system and application data to a remote site. Data is replicated at the byte and file level which provides for a very effective use of available network bandwidth, thus making it ideal for disaster recovery scenarios. As SQL Server is not dependent on Active Directory no domain membership is required, and this further simplifies the fail-over process.

CONFIGURATION

The configuration described here involves setting up CDR to replicate directly to the Standby Server. The Standby SQL Server machine will replace the production machine by assuming its name, but will retain its existing IP address. The original DNS entry for the *Standby Exchange server* will remain intact but the IP address for the *Production Exchange server* will be set to the IP address of the *Standby Server*. The diagram depicts the configuration used in this approach.

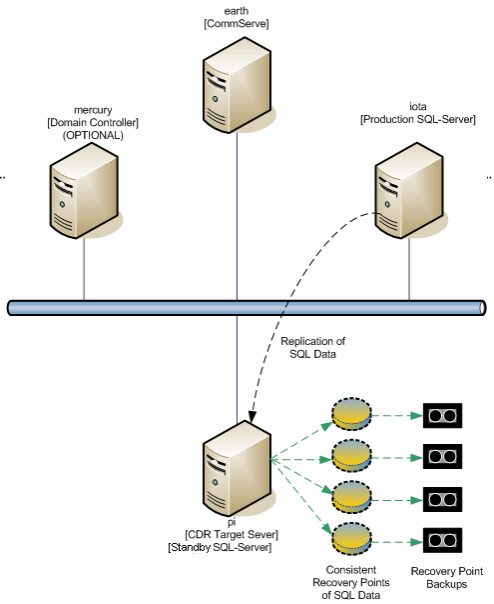

PREPARE THE SQL PRODUCTION SERVER

This procedure assumes that the *Production SQL Server* (iota) is operational and installed with all the required service packs and/or patches. It also assumes that the *Standby SQL Server* (pi), which is also the CDR target, is running as a member of the domain if domain membership is required and that the CommServe has already been installed.

1. Make sure that all databases, including the system databases, are moved off of the C: drive. This is usually the case in production environments but the location should be verified.
2. Install the CDR agent on the server along with the File System agent so the CRPs can be backed up. Considering SQL Server *i*DataAgent is also installed. It is expected that standard backups will be executed on the *Production SQL Server* to manage the logs. Note that CDR will not truncate SQL Server logs since all backup operations associated with CDR are file system based. This means that although a CRP will quiesce SQL Server and provide for a consistent point in time view there is no knowledge of the application during the backup phase. Hence, the SQL Server APIs which truncate logs are not utilized. Note that as the standard SQL Server *i*DataAgent backup causes logs to be truncated on the production server the deletion of the log files will be replicated to the target by CDR, so cleanup will be automatic on the target server.

   NOTE: If using the standard SQL DB agent is not possible on the *Production SQL Server* then the logs must be managed through an alternate mechanism.

PREPARE THE STANDBY SQL SERVER

Configuring the *Standby SQL Server* will require that it be disconnected from the *production server's* network to avoid a name clash. If the *Standby Server* is being created at recovery time then it can be connected to the production network since the original production server will be down. Though these steps could be performed during the actual recovery, it would increase the recovery time and potentially introduce more opportunity for error. So it is recommended that

these steps be performed in advance to provide the most straight forward recovery procedure.

1. Install the ContinuousDataReplicator.
2. Create volumes and folders for storing SQL Server data and logs equivalent to those on the *Production server*. These will be used for replication targets for the live data as well as operational data and logs locations when the *Standby SQL Server* machine is being used.
3. Disconnect the *Standby Server* (pi) from the network.
4. Change the name of the Standby Server (pi) to the name of the production SQL Server (iota) and reboot.
5. Install MS SQL Server using the `/DisasterRecovery`.

   Example: `Z:> setup.exe /DisasterRecovery`

   a. You may receive an informational message on installing the service pack for SQL Server. This situation will be corrected by CommCell and may be disregarded:

   b. When the installation wizard starts make sure that you install all of the components that were on the *Production SQL Server*. Note that you will need to select the same installation folders that were used on the *Production server*.
6. Load the Service Pack(s) and patches that were loaded on the production SQL Server also using the DisasterRecovery switch option The installation wizard will run and install the updates.
7. Make sure that all of the SQL Server services are stopped and set their startup type to Manual.
8. Rename the Standby SQL Server machine back to its original name (pi) and reboot for the change to take effect.
9. Reconnect the Standby SQL Server machine to the production network.

SETTING UP REPLICATION

Now that the servers are prepared and the CommCell software has been, installed the replication set and associated replication pairs can be set up and replication started. The replication set and pairs can be either configured manually, or by the use of wizard. Note the destination paths on the target must be set to mimic the source. See ContinuousDataReplicator for more details on installation and configuration of CDR.

At this point the pairs are in sync and all of the initial data has been replicated to the target (pi). The next steps would be to schedule Consistent Recovery Points along with backups of those Recovery Points. This allows for Point-In-Time recovery for the SQL Server with known consistent data. Alternatively, the live data can be utilized which will provide the closest to up to the minute recovery. However, live data will be crash consistent at best and is not guaranteed to be recoverable by SQL Server.

BRING BACK THE STANDBY SERVER

As there are no tight dependencies with Active Directory recovering the SQL Server application on the Standby SQL machine is rather straight forward. If the machines do belong to a domain then you should DELETE the production SQL Server machine (iota) entry from the Active Directory.

Perform the following steps to bring the *Standby Server* online:

1. It is likely that the initiation of the fail-over to the *Standby SQL Server* machine is due to the *Production SQL Server* crashing but in any case the production server (iota) should be shutdown.
2. If the machines were domain members then on the Domain Controller using Active Directory Users and Computers delete the entry for the *Production SQL Server* (iota).
3. Rename the *Standby SQL Server* (pi) to the *Production SQL Server* (iota) and reboot.
4. In DNS change the IP address of the *Production SQL Server* (iota) to match the IP address of the *Standby SQL Server* (pi). By changing the IP address of the original *Production SQL Server* (iota) to the *Standby Server's* address clients will be able to connect to SQL correctly. Replacing the host entry with an alias record will accomplish the same result.

   Leave the existing entry for the *Standby Exchange server* (pi) intact so that the system can be accessed by its original name as well.
5. At this point the live replica data can be used or a recovery point selected to recover SQL Server. Note that that live data is crash consistent at best and can pre-date the crash point in time due to the asynchronous nature of CDR. However, it will provide the closest to current data for recovery purposes and not require a copy operation since it is in the proper location already. If recovery from the latest replicated data is not desirable or possible due to inconsistencies, then a CRP or backup can be chosen instead. From the CommCell Console browse the CRPs and select the one to be utilized on the *Standby Server*.
6. Recovery of CRP data can be done two ways. (Use one method.)

   a. The Copyback button can be used to initiate the copy operation from the GUI. The original *Standby SQL Server* (pi) is selected as the recovery host and the Copyback function is initiated. This is the most straight forward method.

      NOTE: The copy back will overwrite the live copy so be sure that the live copy will not be required in the future. Otherwise manually copy the data to a new location before issuing the Copyback function. After selecting the Standby SQL Server machine by its original name (pi) click OK to start the copy operation.

b. Alternatively, the CRP could be mounted and manually copied to the required location. This is accomplished through the **View Snapshots** function. Click View Snapshots Each of the snapshots in the CRP will be listed and each should be mounted by selecting the individual snapshot and clicking the Mount button which exposes the mount point dialog so that a drive letter can be assigned.

7. Start the SQL services on the *Standby SQL Server* and set the Startup Type for each to automatic if the configuration will remain intact for a sufficient period of time.
8. Verify that SQL is operational and all databases are accessible.

Back to Top