Features - Image Level *i*DataAgent

TABLE OF CONTENTS

Overview - Image Level *i*DataAgent

Choose from the following topics:

- Introduction
- Supported Data Types
- Tree Levels in the Image Level *i*DataAgent
- License Requirement
- Snapshot Support
- LAN Copy Manager
- Disaster Recovery Considerations

INTRODUCTION

Unlike the file system *i*DataAgent which backs up folders and files, the Image Level *i*DataAgent backs up complete volumes or mount points, and restores them using Volume Level Restore; for certain operating systems it is also able to restore folders and files using a File Level Restore.

Generally, only blocks that contain data are backed up; empty blocks are not, with the following notes:

- For FAT16/FAT32 volumes, all blocks on a drive are backed up, even if the blocks contain no data.
- For Image Level on Unix, this functionality can be turned off, so that even the empty blocks are backed up, by configuring the `dDisableDataBlock` registry key.
- CXBF devices on AIX do not support online file system expansion or shrink using the `chfs` command. If a volume or file system controlled by a CXBF block-filter driver is extended or shrunk, it will fail. Therefore, you must de-configure the CXBF device before using `chfs`.

Server downtime is almost completely eliminated by capturing an image of a production server's disk with data snapshot technology (e.g., QSnap), then using the Image Level *i*DataAgent to back up the snapshot.

SUPPORTED DATA TYPES

The following file system types are supported for backup and restore operations:

IMAGE LEVEL ON WINDOWS

- File Allocation Table (FAT) file systems - Volume Level Restore only
- New Technology File Systems (NTFS) - both Volume Level and File Level Restores
- RAW volumes - Volume Level Restores only
- VxVM Volumes (VxVM 3.1 and higher)

For Windows data, it is possible to restore data from one file system type to another through Volume Level Restore. The following such Volume Level Restores are supported:

- NTFS ==>> FAT
- NTFS ==>> RAW
- FAT ==>> RAW

After such a restore, sometimes Windows Explorer may still show the volume having the same file system type as before the restore, even though Computer Management shows the new (correct) type. After a reboot, Windows Explorer will show the correct file system type. For File Level Restores, only data backed up from NTFS volumes can be browsed and restored to NTFS or FAT volumes.

- For Windows, when volumes are backed up using Image Level, sharing attributes for volumes and folders are not retained.
- For Windows, File Level restores require the MediaAgent Index Cache to reside on an NTFS partition. In the case of a MediaAgent whose Index Cache resides on a FAT partition, use Volume Level restores.

IMAGE LEVEL ON UNIX

The following table lists the file systems and restore types supported for each supported operating system:

File System	Operating System	CXBF	Checksum	Back Up Data Blocks Only		Volume Restore	File Restore
				CXBF	Checksum		

Unix File System (UFS) **Includes read-only partitions on Unix, as long as the production server and backup host are configured as the same computer**	Solaris	X		X		X	X
Extent 2 File System (ext2)	Linux	X	X	X		X	X
Extent 3 File System (ext3)	Linux	X	X	X		X	X
RAW volumes	Linux Solaris	X				X	
Reiser File System (reiserfs)	Linux	X	X			X	
VERITAS File System (VxFS)	AIX		X			X	
VERITAS File System (VxFS)	Solaris	X				X	
Veritas Volume Manager (VxVM)	Solaris	X	X	X		X	
Journal File System (JFS2)	AIX	X	X	X	X	X	X
'X' File System (XFS)	Linux	X	X			X	

- The Unix File System snapshots cannot be backed up by non-CXBF backup jobs.
- The Unix QSnap® (cxbf) driver does **not** support volumes created by Solaris Volume Manager.
- On AIX clients, Data Blocks Only backups are supported for full backups only.
- The Image Level and Image Level ProxyHost *i*DataAgents on Unix can perform a File Level Restore only when the OS of the Client computer and the MediaAgent are the same.

For a complete listing of applications supported for each operating system, see Image Level - Application Support.

Back to Top

TREE LEVELS IN THE IMAGE LEVEL *i*DATAAGENT

When the Image Level *i*DataAgent is installed, the following levels are automatically created in the CommCell Browser:

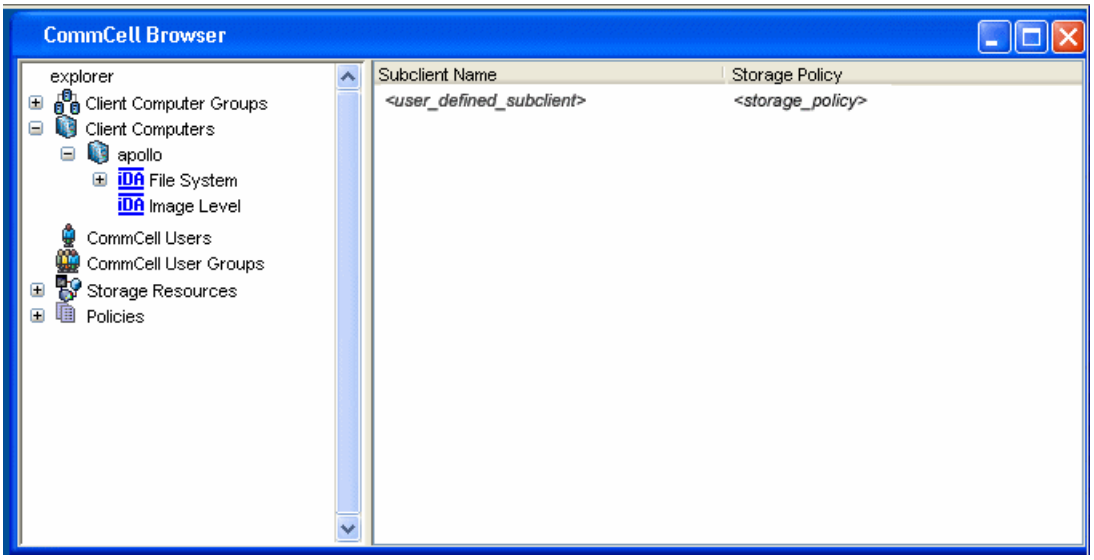

`apollo`: Client

`<user_defined_subclient>`: User-defined subclients

`Image Level`: Agent

Back to Top

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

SNAPSHOT SUPPORT

The Image Level *i*DataAgent is designed to work in conjunction with the following snapshot engines, which provide snapshot functionality for data protection operations:

- VSS for the Image Level *i*DataAgent
- QSnap for the Image Level *i*DataAgent

For more detailed snapshot support information, refer to Support Information - Snapshot Engines.

Back to Top

LAN COPY MANAGER

The Image Level *i*DataAgent uses LAN Copy Manager to read block-level data directly from the specified volume(s) and transmit it to the designated MediaAgent over the LAN, which then writes the received data directly to the destination. For more information, see Copy Managers.

Back to Top

DISASTER RECOVERY CONSIDERATIONS

- Before you use your agent, be sure to review and understand the associated full system restore (or disaster recovery) procedure. The procedure for some agents may require that you plan specific actions or consider certain items before an emergency occurs. See Disaster Recovery for more information regarding your agent.
- Image Level *i*DataAgent cannot be used for disaster recovery of operating system volumes. For agents that support disaster recovery, refer Disaster Recovery.

Back to Top

System Requirements - Image Level *i*DataAgent

System Requirements | Snapshot Support | Supported Features

The following requirements are for the Image Level *i*DataAgent:

If using this Agent with a snapshot engine, refer to System Requirements-Image Level *i*DataAgent - Snapshots Support information on the operating system vendors supported for each snapshot engine with this Agent.

OPERATING SYSTEM		PROCESSOR
AIX	AIX 7.1 64-bit	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	Power PC (Includes IBM System p)
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
	Debian 4.x with kernel 2.6.18-686	Intel Pentium or compatible minimum required
	Debian 4.x with kernel 2.6.18-6	x64
	RED HAT ENTERPRISE LINUX	
	Red Hat Enterprise Linux AS 4.0 with kernel 2.6.9-89 (Update 8)	Intel Pentium or compatible minimum required
	Red Hat Enterprise Linux AS 4.0 with kernel 2.6.9-89 (Update 8)	x64
	Red Hat Enterprise Linux AS 4.0 with kernel 2.6.9-78 (Update 7)	x64
	Red Hat Enterprise Linux AS 4.0 with kernel 2.6.9-78 (Update 7)	Intel Pentium or compatible minimum required
	Red Hat Enterprise Linux AS 4.0 with kernel 2.6.9-67 (Update 6)	Intel Pentium or compatible minimum required
	Red Hat Enterprise Linux AS 4.0 with kernel 2.6.9-67 (Update 6)	x64
	Red Hat Enterprise Linux AS 4.0 with kernel 2.6.9-42 (Update 4)	x64
	Red Hat Enterprise Linux AS 4.0 with kernel 2.6.9-42 (Update 4)	Intel Pentium or compatible minimum required
	Red Hat Enterprise Linux AS 4.0 with kernel 2.6.9-34 (Update 3)	Intel Pentium or compatible minimum required
	Red Hat Enterprise Linux AS 4.0 with kernel 2.6.9-34 (Update 3)	x64
	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 2.6.32-279 (Update 3)	x64
	Red Hat Enterprise Linux 6 Advanced Platform with kernel 2.6.32-279 (Update 3)	Intel Pentium or compatible minimum required
	Red Hat Enterprise Linux 6 Advanced Platform with kernel 2.6.32-220 (Update 2)	x64
	Red Hat Enterprise Linux 6 Advanced Platform with kernel 2.6.32-220 (Update 2)	Intel Pentium or compatible minimum required
	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 2.6.18-308 (Update 8)	x64
	Red Hat Enterprise Linux 5 Advanced Platform with kernel 2.6.18-308 (Update	Intel Pentium or compatible minimum required

	8)	
	Red Hat Enterprise Linux 5 Advanced Platform with kernel 2.6.18-274 (Update 7)	Intel Pentium or compatible minimum required
	Red Hat Enterprise Linux 5 Advanced Platform with kernel 2.6.18-274 (Update 7)	x64
	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 2.6.18-238 (Update 6)	Intel Pentium or compatible minimum required
	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 2.6.18-194 (Update 5)	x64
	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 2.6.18-164 (Update 4)	Intel Pentium or compatible minimum required
	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 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)	x64
	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 2.6.27.19-5	x64
	SuSE Linux 11 Enterprise Server with kernel 2.6.27.19-5	Intel Pentium or compatible minimum required
	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 2.6.16.60-0.54.5 (Update 3)	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 2.6.16.60-0.21 (Update 2)	x64
	SuSE Linux 10 Enterprise Server with kernel 2.6.16.46-0.12 (Update 1)	x64
	SuSE Linux 10 Enterprise Server with kernel 2.6.16.46-0.12 (Update 1)	Intel Pentium or compatible minimum required
SOLARIS	Solaris 9 with a minimum of Service Packs 111711-02	Sparc5 or higher recommended
	Solaris 10.x with a minimum of SunOS (Sparc) Patch 119963-14	Sparc5 or higher recommended
WINDOWS	**WINDOWS 2008**	
	Microsoft Windows Server 2008 32-bit and x64 Editions* *Core Editions not supported	All Windows-compatible processors supported
	WINDOWS VISTA	
	Microsoft Windows Vista 32-bit and x64 Editions	All Windows-compatible processors supported
	WINDOWS XP	
	Microsoft Windows XP Professional Edition 32-bit with a minimum of Service Pack 3	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

WINDOWS

115 MB minimum of hard drive space for software/ 499 MB recommended

100 MB of additional hard disk space for log file growth

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

UNIX

235 MB minimum of hard drive space for software

MEMORY

WINDOWS

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

AIX, LINUX, AND SOLARIS

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

Swap space = 2*RAM size

AIX LPAR/WPAR SUPPORT

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

PERIPHERALS

DVD-ROM drive

Network Interface Card

MISCELLANEOUS

The File System *i*DataAgent will be automatically installed during installation of this software, if it is not already installed. For System Requirements and install information specific to the File System *i*DataAgents, refer to:

- System Requirements - Microsoft Windows File System *i*DataAgent
- System Requirements - AIX File System *i*DataAgent
- System Requirements - Linux File System *i*DataAgent
- System Requirements - Solaris File System *i*DataAgent

On Solaris computers, the operating system must have been installed with at least the `user level software` option selected.

NETWORK

TCP/IP Services configured on the computer.

.NET FRAMEWORK

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

MICROSOFT VISUAL C++

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

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 backup_IDA application, you can use the following file name: backup_IDA.fc	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 (system_u:object_r:texrel_shlib_t,s0) /opt/<software installation directory>/Base/libdb2osse.so.1 -- gen_context (system_u:object_r:texrel_shlib_t,s0) /opt/<software installation directory>/Base/libDb2Sbt.so -- gen_context (system_u:object_r:texrel_shlib_t,s0) /opt/<software installation directory>/Base/libdb2trcapi.so.1 -- gen_context (system_u:object_r:texrel_shlib_t,s0) /opt/<software installation directory>/Base/libDrDatabase.so -- gen_context (system_u:object_r:texrel_shlib_t,s0) /opt/<software installation directory>/Base/libIndexing.so -- gen_context (system_u:object_r:texrel_shlib_t,s0) /opt/<software installation directory>/Base/libSnooper.so -- gen_context (system_u:object_r:texrel_shlib_t,s0)

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

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

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

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

SELinux is now configured to work with this application.

DISCLAIMER

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

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

Install the Image Level *i*DataAgent - Windows

TABLE OF CONTENTS

INSTALL REQUIREMENTS

The following procedure describes the steps involved in installing the Image Level *i*DataAgent and the QSnap snapshot enabler on both cluster and non-clustered environment.

The Image Level *i*DataAgent is installed on the computer from which the *i*DataAgent secures data. (This computer is referred to as the *Client* computer in this install procedure.)

Verify that the computer in which you wish to install the software satisfies the minimum requirements specified in System Requirements - Image Level *i*DataAgent and System Requirements - Microsoft Windows File System *i*DataAgent.

Review the following Install Requirements before installing the software:

GENERAL

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

CLUSTER SPECIFIC

- Before you can install this Agent in the cluster group in a clustered environment, you must first install QSnap on all physical nodes of the cluster. Install QSnap - Windows provide step-by-step instructions for installing these components on physical nodes.
- Once the Windows File System *i*DataAgent and QSnap are installed on all physical nodes in the cluster, the Image Level *i*DataAgent and Windows File System *i*DataAgent 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.
- Check the following on the cluster computer in which you wish to install the software:
  - Cluster software is installed and running.
  - Active and passive nodes are available.
  - Disk array devices configured with access to the shared array.
  - Public Network Interface Card is bound first, before the private Network Interface Card. (Does not apply to NetWare Cluster.)

BEFORE YOU BEGIN

- Log on to the client as 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.

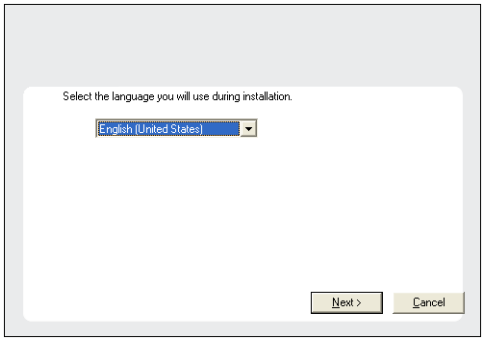

3. Select the option to install software on this computer.

   NOTES

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

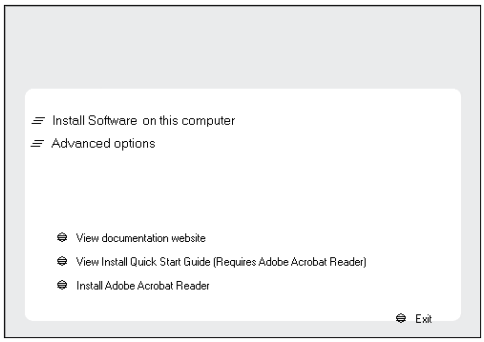

4. Read the license agreement, then select **I accept the terms in the license agreement**.

   Click **Next** to continue.

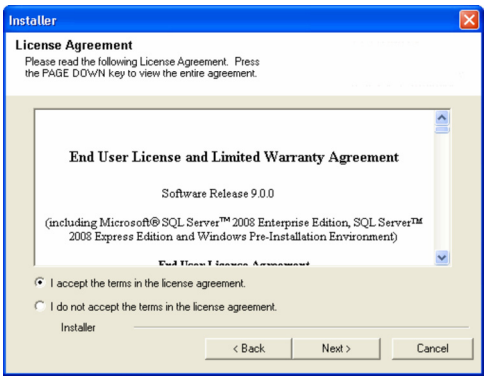

CLUSTER SELECTION

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

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

   Click **Next** to continue.

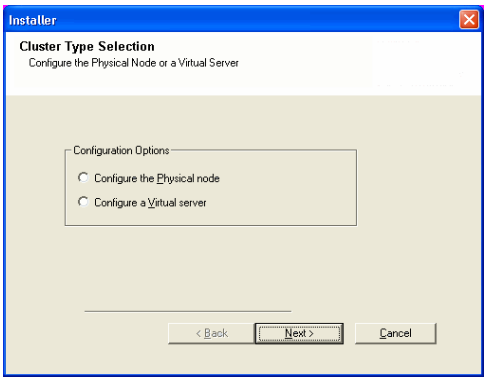

6. Select the disk group in which the cluster group resides.

   Click **Next** to continue.

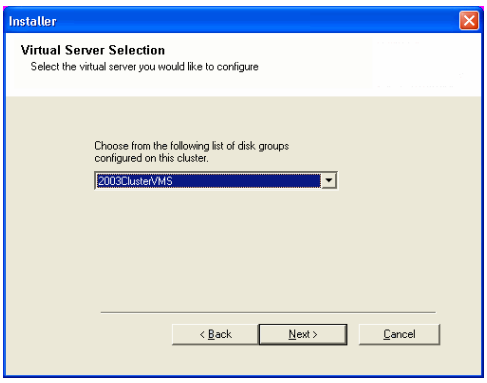

SELECT COMPONENTS FOR INSTALLATION

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.

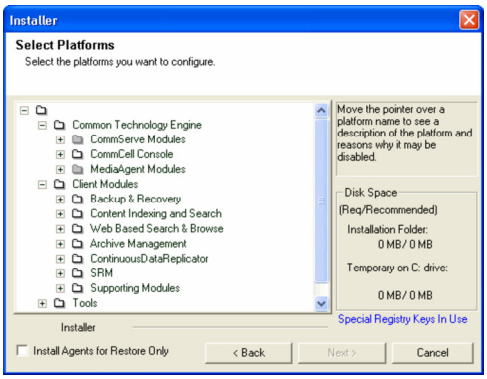

   To install the Image Level *i*DataAgent, expand the following `Client Modules` folder, `Backup & Recovery` folder and `File System` folder. Then select the following:

   - `Image Level iDataAgent`

   When you select the Image Level *i*DataAgent for installation, the appropriate Windows File System *i*DataAgent and the QSnap snapshot enabler are automatically selected for installation.

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.

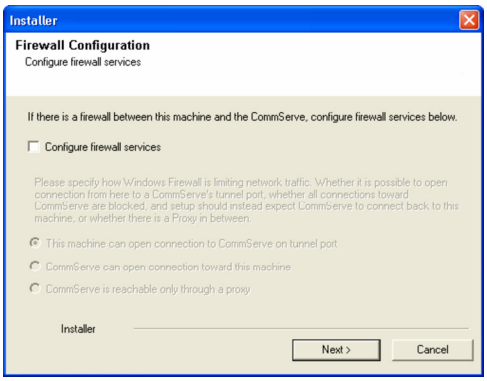

9. Enter the fully qualified domain name of the CommServe Host Name. This should be TCP/IP network name. e.g., computer.company.com.

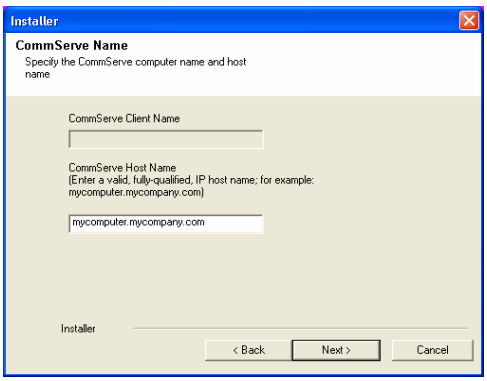

NOTES

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

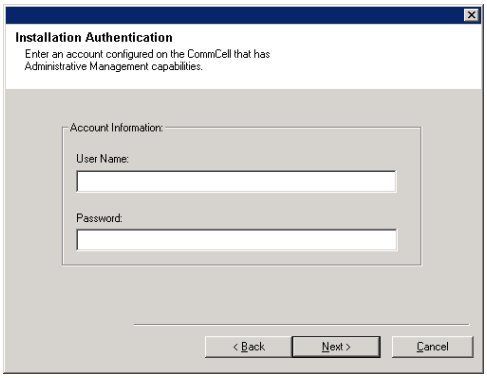

NOTES

- 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.

11. 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.

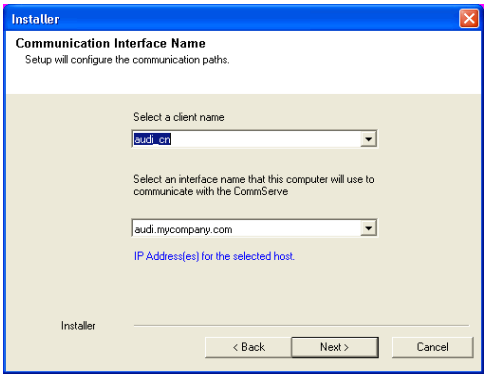

NOTES

- 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.

12. 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.

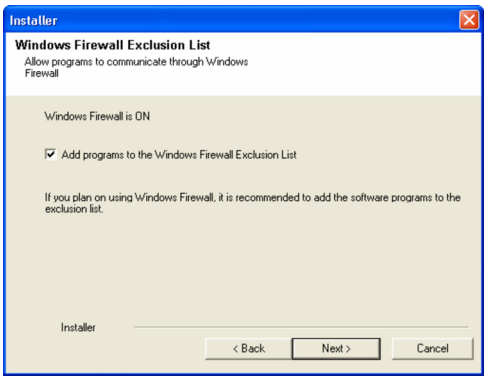

13. Specify the location where you want to install the software.

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.

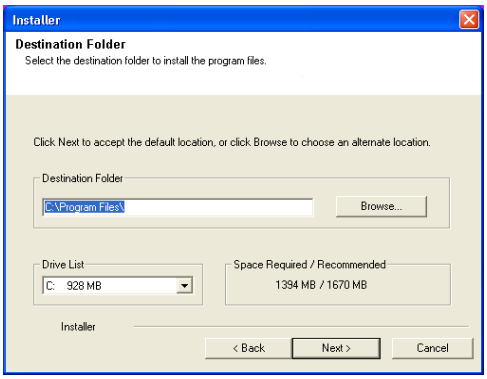

CLIENT GROUP SELECTION

14. 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.

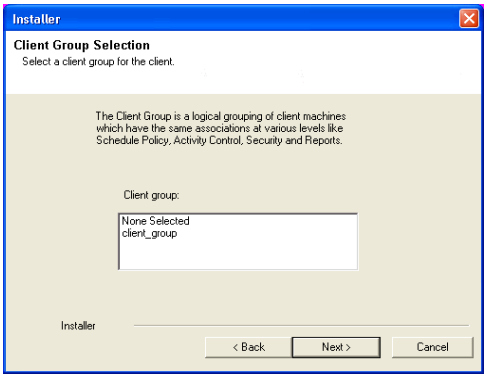

SCHEDULE AUTOMATIC UPDATE

15. 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.

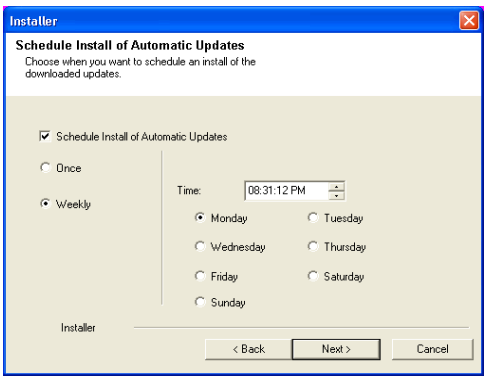

STORAGE POLICY SELECTION

16. Select the storage policy through which you want to back up/archive the agent.

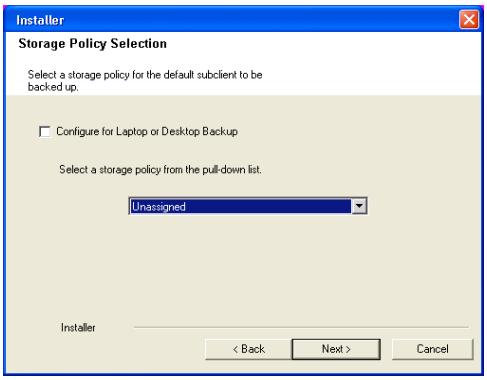

NOTES

- A storage policy directs backup data to a media library.
- If desired, you can change your storage policy selection at any time after you have installed the client software.
- This screen may appear more than once, if you have selected multiple agents for installation. You will be prompted to configure the storage policy association for each of the selected agents.

Click **Next** to continue.

CONFIGURE QSNAP

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

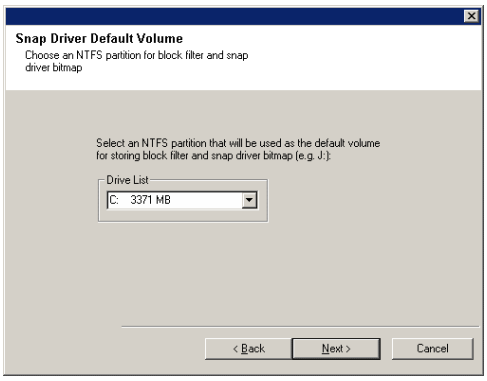

NOTES

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

Click **Next** to continue.

18. Verify the summary of selected options.

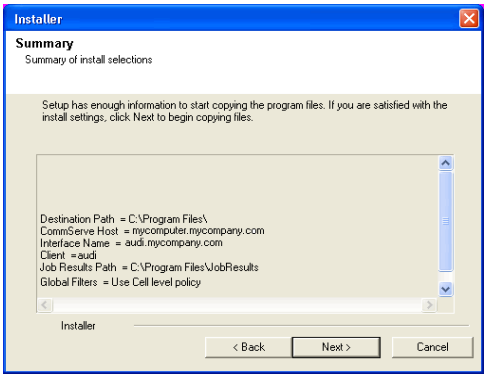

NOTES

- 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.

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

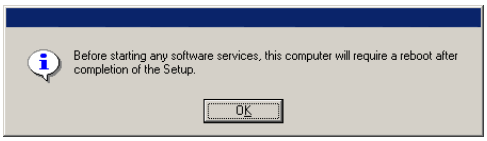

Click **OK** to continue.

VERIFY SUMMARY OF INSTALL OPTIONS

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

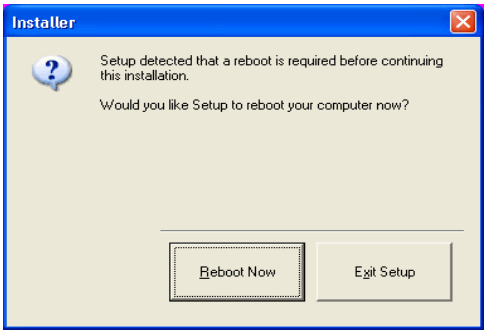

- **Reboot Now**
  If this option is displayed without the **Skip Reboot** option, the install program has found files required by the software that are in use and need to be replaced. If **Reboot Now** is displayed without the **Skip Reboot** option, reboot the computer at this point. The install program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

INSTALL REMAINING CLUSTER NODES

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

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

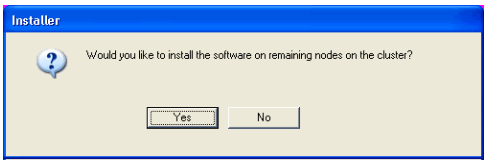

22. Select cluster nodes from the **Preferred Nodes** list and click the arrow button to move them to the **Selected Nodes** list.

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.

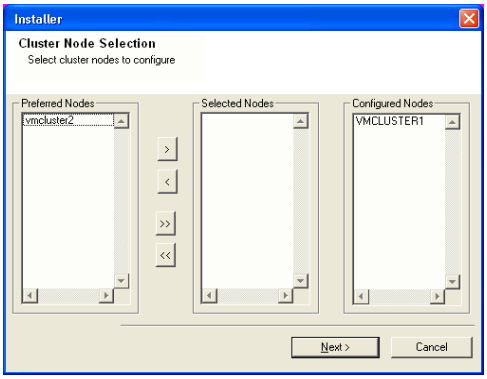

23. Type the **User Name** and **Password** for the Domain Administrator account, so that the installer can perform the remote install/upgrade of the cluster nodes you selected in the previous step.

Click **Next** to continue.

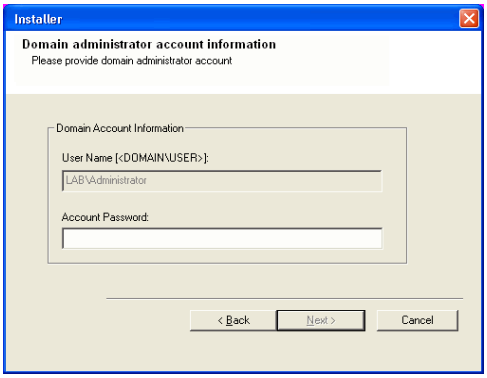

24. The progress of the remote install for the cluster nodes is displayed; the install can be interrupted if necessary.

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.

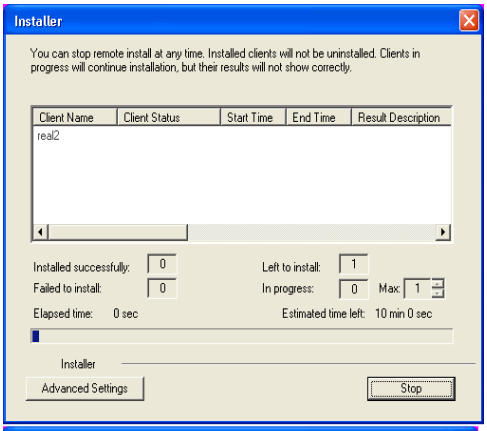

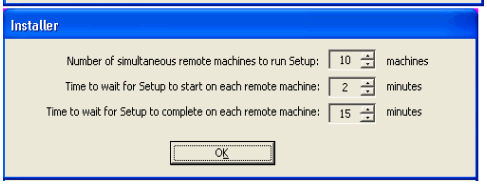

25. Read the summary for remote installation to verify that all selected nodes were installed successfully.

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.

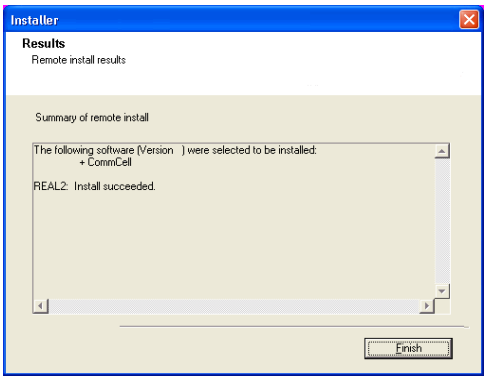

SETUP COMPLETE

26. 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.

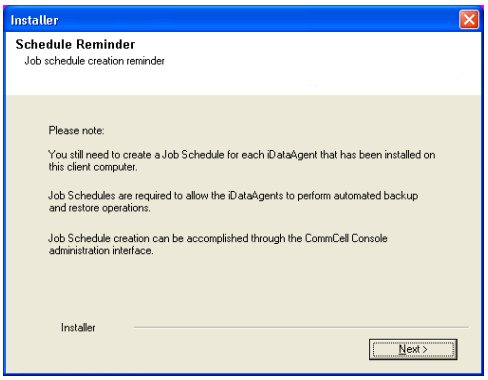

27. Select from the following:

- 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.

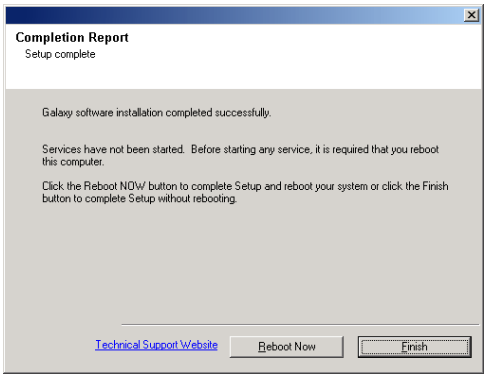

POST-INSTALL CONSIDERATIONS

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 Image Level *i*DataAgent - UNIX

TABLE OF CONTENTS

INSTALL REQUIREMENTS

The following procedure describes the steps involved in installing the Image Level and Unix File System *i*DataAgents, and QSnap (CVGxCXBF). The Image Level *i*DataAgent is installed on the computer from which the *i*DataAgent secures data. (This computer is referred to as the *Client* computer in this install procedure.)

Verify that the computer in which you wish to install the software satisfies the minimum requirements specified in the following:

- System Requirements - Image Level *i*DataAgent
- System Requirements - AIX File System *i*DataAgent
- System Requirements - HP-UX File System *i*DataAgent
- System Requirements - Linux File System *i*DataAgent
- System Requirements - Solaris File System *i*DataAgent

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.

RED HAT LINUX

Red Hat Linux will create an entry in the `/etc/hosts` file when it is first installed, in the following format:

<ip_address> *<host name>* `localhost`

For example, if the host name of your computer is `bluesky`, the entry will look something like this:

```
192.168.1.111 bluesky localhost
```

If you have not already done so, edit the `/etc/hosts` file. The edited entry should look like this:

```
127.0.0.1 localhost
```

Depending upon your environment, and using the above example again, you may also need an entry similar to this:

```
192.168.1.111 bluesky
```

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.

   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.

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

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

1) Install data protection agents on this computer

2) Advance options

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 **CVGxImgIDA** 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.

```
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 ?
```

- 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.)

```
=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** to continue. See Installing Restore Only Agents for more information.

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

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

19. 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 ...
```

KERNEL PARAMETERS

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

    NOTES

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

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

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

Number of streams: [10]
```

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

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

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

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

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

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

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

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

```
Press <ENTER> to continue.

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

set shmsys:shminfo_shmmax=4199304
set shmsys:shminfo_shmmin=1
set semsys:shminfo_shmmni=640
set semsys:shminfo_shmseg=640
set semsys:seminfo_semmns=640
set semsys:seminfo_semmni=640
set semsys:seminfo_semmsl=640
set maxusers=256
```

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

NOTES

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

Press **Enter** to continue.

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

ENABLE GLOBAL FILTERS

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

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

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

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

NOTES

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

CLIENT GROUP SELECTION

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

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

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

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

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

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

NOTES

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

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

IMPORTANT:

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

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

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

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

25. Press **Enter** to continue.

NOTES

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

STORAGE POLICY SELECTION

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

1) SP_StandAloneLibrary2_2
2) SP_Library3_3
3) SP_MagLibrary4_4
4) fornax_fornax_HWCmp
5) ranger_ranger_HWCmp
6) fornax_fornax_ClntCmp
7) fornax_fornax_MACmp
8) fornax_fornax_NoCmp

Storage Policy: [3]
```

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

NOTES

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

agents for installation and are configuring storage policy association for each of the installed agents.

BLOCK FILTER AND SNAP DRIVER (CVGXCXBF) INSTALLATION

27. To install **CVGxCXBF** package, type **Yes** and press **Enter** to continue.

NOTES:

- Skip to Setup Complete to install this package at a later time.

```
Cxbf provides snapshot facility and also monitors the
device to provide incremental support. You can chose to
install now or you can install later.

Do you want to install? [Yes]
```

28. Press **Enter** to continue.

```
For every snapshot taken CXBF driver creates a sparse file
where it puts copy-on-write data blocks and block bitmap.
These COW files are created in a dedicated directory that
should satisfy a few requirements:

1) It should be a local directory for performance reasons,

2) It should have sufficient amount of space,

3) In case of Solaris UFS, it should have journaling
turned off.

COW Cache Directory: [/space/opt/calypso2/CXBF/cache]
```

29. Make your selection and press **Enter**.

NOTES

- This prompt is displayed only if the directory identified in the previous Step does not exist.

```
Directory "<dir_name>" does not exist.

1) Create this directory and continue
2) Retry input

Your choice: [1]
```

30. Record this mount point for use later when mounting a volume to it.

Press **Enter**.

```
Please make a note of this mount point and use it for
mounting the cache volume in the Volume Explorer after
install has completed.

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

31. The install program now starts copying the software to the computer. The progress of the operation is displayed.

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

Successfully copied xx files
```

32. To grant this permission (recommended), type **yes**. Otherwise, edit the `/etc/rc0` system script as shown and type **no**.

NOTES

- If you do not want to grant this permission, you must edit the `/etc/rc0` system script now, by inserting the following commands after the line that unmounts all application volumes but before the line that unmounts system partitions:

  ```
  echo "Bull CXBF: updating the cxbf on disk bitmaps."

  /usr/sbin/cvsnap -c "updatebitmaps"
  ```

```
We now need to make some changes to your /etc/rc0 system
script.

We can either make these changes ourselves at this time,
or we can explain what has to be modified so that you can
do it yourself.

Would you like us to modify /etc/rc0 now? [yes]
```

33. Press **Enter** to continue.

```
Successfully installed CVGxCXBF.

Press ENTER to proceed with CVGxImgIDA install ...
```

34. The install program now starts copying the software to the computer. The progress of the operation is displayed.

Press **Enter** to continue.

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

Successfully copied xx files

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

Successfully installed CVGxImgIDA

Press ENTER to continue ...
```

SETUP COMPLETE

35. Enter the number corresponding to the **Exit** option and then press **Enter** to continue.

The installation is now complete.

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

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

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

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

Your choice: [1]
```

POST-INSTALL CONSIDERATIONS

GENERAL

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

Backup - Image Level

Topics | How To | Related Topics

OVERVIEW

Plan your backup jobs for this agent by reviewing the following information:

- For an overview of backup jobs, see Backup Data.
- For a list of supported data types for this agent, see Supported Data Types.
- For information on subclients, see Subclients
  - For information on configuring subclients for this agent, see Subclients - SAN *i*DataAgents.

SUPPORTED BACKUP TYPES

This agent supports the following backup types:

- Full Backups
- Incremental Backups
- Synthetic Full Backups
- CXBF Backups
- Non-CXBF (Checksum) Backups
- Oracle BLI

IMAGE LEVEL BACKUP JOBS

This agent has the following unique functionality and options for protecting data:

The Image Level *i*DataAgent backs up the blocks of complete volumes, and restores them using Volume Level Restore or File Level Restore. Each backup for the *i*DataAgent on a Unix platform is either a CXBF backup or a non-CXBF backup.

CXBF BACKUPS

CXBF backups are run on Unix platforms for this agent. To run a CXBF backup, QSnap must be installed on the client and you must configure a CXBF subclient. For more information, see Subclients - SAN *i*DataAgents. For a step-by-step procedure, see Run a CXBF Backup (Image Level/Image Level ProxyHost).

NON-CXBF BACKUPS

Non-CXBF (Checksum) backups are run on Unix platforms for this agent. To run a non-CXBF (Checksum) backup, you must configure a non-CXBF (Checksum) subclient and also provide a supported snapshot. For more information, see Subclients - SAN *i*DataAgents. For a step-by-step procedure, see Run a non-CXBF Backup (Image Level/Image Level ProxyHost).

BACKING UP APPLICATION DATA

To facilitate the backup of data files for application such as Exchange and SQL, PreScan and PostScan commands can be utilized to quiesce and unquiesce the application during the backup. For more information, see Pre/Post process commands.

ORACLE BLI

To be able to back up Oracle instances with the Image Level *i*DataAgent, the Oracle data files and Oracle archive logs must reside on two separate volumes. In addition, the Oracle installation binaries, control file, and online redo logs must reside in volumes other than the data files volume and the archive logs volume.

During an Oracle backup, each Image Level subclient is expected to quiesce one instance at a time. Do not configure an Image Level subclient's content with more than one Oracle instance; create separate subclients for each Oracle instance.

Pre/Post process commands are required to quiesce and unquiesce the Oracle instance during the backup. See the Agent-Specific Guidelines for more information.

For step-by-step instructions, see Back up Oracle Instances with the Image Level *i*DataAgent.

CONFIGURING MULTIPLE STREAMS FOR BACKUPS

Multi-Streaming employs multiple data streams per subclient for data protection operations. This enables the subclient's contents to be distributed to all the streams, transmitting them in parallel to the storage media. Hence a subclient whose data is secured using three data streams, utilizes more of the available network resources, and can complete in as little as one third the time that the same data would require using a single stream.

Note that, you do not have to enable multi-streaming if you have multiple mount points pointing to the same physical drive.

For Unix, multiple streams are supported for CXBF backups. Follow the steps given below to configure multi-streaming.

1. From the CommCell Browser, navigate to **Client Computers | <*Client*> | Image Level on Unix**.
2. Right-click the **<*Subclient*>** in the right pane and then click **Properties**.
3. Click **CXBF**.
4. In the **Number of Data Readers** box type or select the number of data streams.
5. Select the **Allow multiple data readers within a drive or mount point** check box.
6. Click the **Storage Device** tab.
7. In the **Storage Policy** list, click a storage policy name.
8. Click **OK**.

The number of streams configured in the Storage Policy should be equal to, or greater than the specified **Number of Data Readers** .

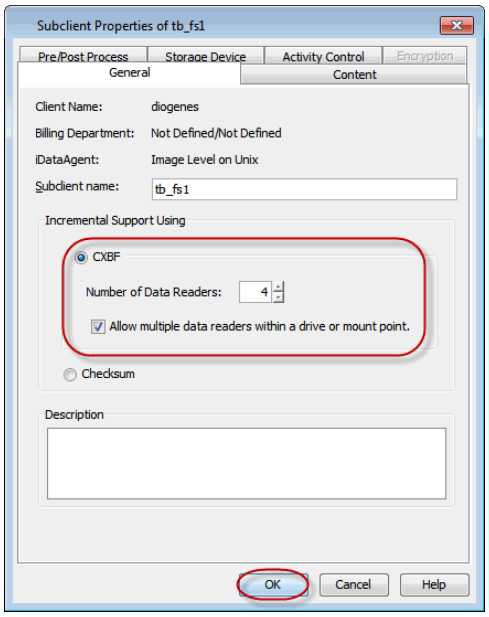

ENHANCING BACKUP PERFORMANCE ON AIX CLIENTS

Image level backup performance on AIX clients can be increased by increasing the amount of data read (chunk size) in one attempt during backups. By default, the chunk size is 4096 Bytes. For a typical read size of 61440 Bytes, it will take 15 data read iterations to complete the data fetch. If the chunk size is increased to 65536 Bytes, only one data read iteration is needed.

INCREASING THE DATA READ SIZE

Use the following steps to increase the amount of data read (chunk size) in one attempt:

1. From the CommCell Console, navigate to **Client Computers**.
2. Right-click the <***Client***>, and then click **Properties**.
3. Click the **Registry Key Settings** tab.
4. Click **Add**.
5. In the **Name** field, type `dSnapChunkSize`.

6. In the **Location** list, type `UnixImageIDA`.
7. In the **Type** list, select `Value`.
8. In the **Value** field, type `65536`.
9. Click **OK**.

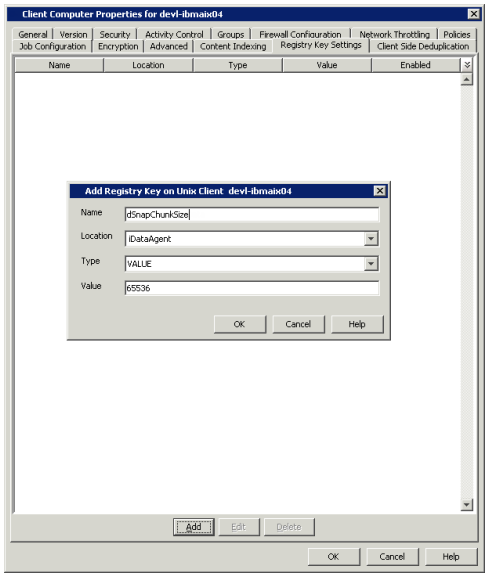

ENABLING DATA CLASSIFICATION FOR METADATA COLLECTION

Metadata contains information such as the directory structure and file attributes. Metadata collection at the file level during backups is necessary to perform file level restores. Metadata collection performance can be enhanced using Data Classification. When Data Classification is used for metadata collection, ACLs will not be available for restore

On a Unix platform, a CXBF device is a volume or partition that is monitored by the CXBF block-filter driver. Consider the following when configuring a CXBF device:

- CXBF devices should not be configured on Operating System file systems.
- CXBF devices can only be configured on quiescent file systems (i.e. unmountable or not busy).

1. Ensure that the Data Classification database is located on a non Data Classification monitored volume and a non CXBF monitored volume.
2. At the Unix prompt on the client computer, type the command to define the location of the Data Classification database after navigating to the `/opt/Calypso/Base` location:

   ```
   DcClient -edit DB_FOLDER <dbpath>
   ```

   The `<dbpath>` must exist prior to running the command.
3. From the CommCell Browser, navigate to **Client Computers | *<Client>* | Image Level on Unix**.
4. Right-click the ***<Subclient>*** and click **Properties**.
5. Click **Volume Explorer**, and then click **Yes** on the warning dialog box.
6. Select the host connected to the volume you want to configure, and then right-click the volume and select **Configure CXBF device**.
7. Click **Configure** in the **Submit Configure Request** dialog box.
8. Click **OK**.
9. At the Unix prompt on the client computer, type the command to start monitoring the CXBF mount point after navigating to the `/opt/Calypso/Base` location:

   ```
   DcClient -monitor <mount path>
   ```

   where: `<mount path>` is the CXBF mount point that you want to add to the monitoring list.
10. From the CommCell Console, navigate to **Client Computers**.
11. Right-click the <***Client***>, and then click **Properties**.
12. Click the **Registry Key Settings** tab.
13. Click **Add**.
14. In the **Name** field, type `sUseDCForMetaData`.
15. In the **Location** list, type `UnixImageIDA`.
16. Click **OK**.

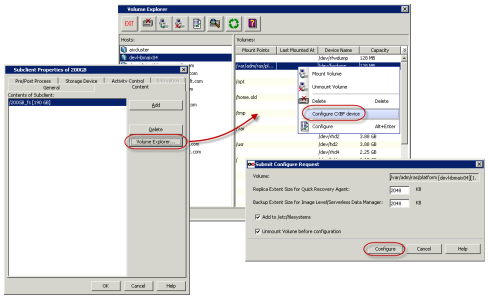

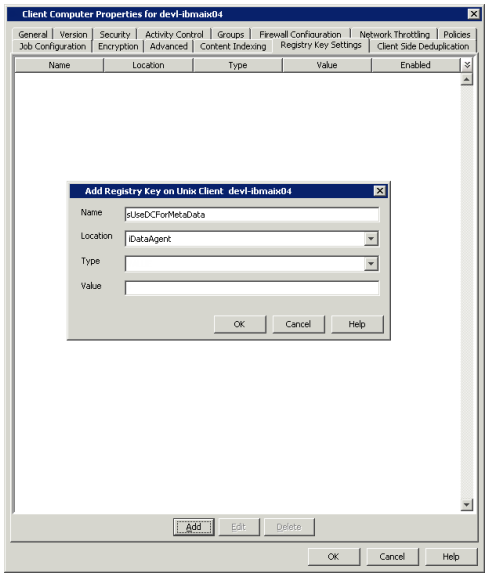

DISABLING RAW DEVICE BACKUPS ON AIX CLIENTS

Raw device backups are enabled by default on AIX clients. However, you can disable raw device backups and allow block device based backups using the following steps:

1. From the CommCell Console, navigate to **Client Computers**.
2. Right-click the <***Client***>, and then click **Properties**.
3. Click the **Registry Key Settings** tab.
4. Click **Add**.
5. In the **Name** field, type `sSnapUseBlkDevice`.
6. In the **Location** list, type `UnixImageIDA`.
7. In the **Type** list, select `Value`**.**
8. In the **Value** field, type `Y`.
9. Click **OK**.

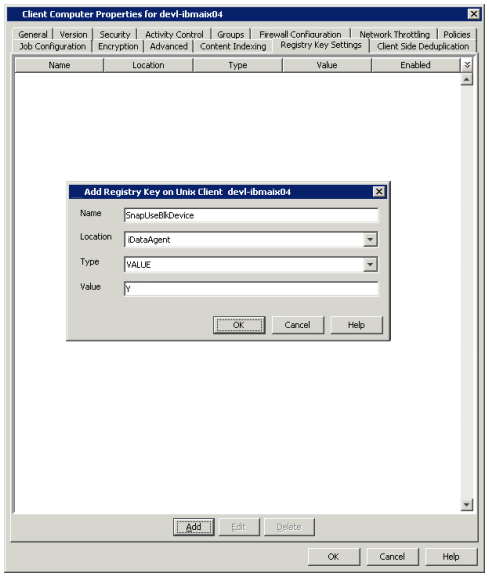

BACKUP CONSIDERATIONS

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

- **Metafile Creation**
  - Selecting the Skip Metafile Creation option in the Advanced Backup Options screen will increase performance but disables the ability to restore data using a File Level restore. This will also apply to any Synthetic Full backups created from such a backup.
  - If File Level restores will never be used for this client, instead of manually selecting the Skip Metafile Creation option for each job, edit the `SkipMetaFileCreation` registry key to automatically skip metafile creation for all backup jobs.
  - For the Image Level and Image Level ProxyHost *i*DataAgents, if a backup job is suspended either by the user or the Job Controller during metadata collection, the job will automatically resume from the scan phase.
- For increased logging of activities during data protection operations, the `dEnableIBackupLog` registry key can be created.
- To alter the default snapshot chunk size, the `dSnapChunkSize` registry key can be created.
- If the backup extent size is changed, the next backup will be converted to a full backup.
- For the Windows Image Level *i*DataAgent, you can add volumes to existing subclient content or remove them. However, whenever volumes are added or removed, the next backup job will be converted to a full backup.
- The cluster size (allocation unit) on any disk that you back up must be greater than 1024 bytes if you intend to restore the data using a file level restore. See Restore Data - Image Level - File Level Restore.
- For Solaris, under conditions of heavy I/O, such as is expected during block-level backups and QR volume creation, we recommend that you do not enable the UFS logging option on the client machine.
- For Windows, before performing the first full backup of a volume, defragment it. This will ensure that the minimum number of extents will need to be backed up, resulting in better performance and lower storage requirements.
- Although you can use the Image Level *i*DataAgent to back up a virtual server's volumes, you cannot use the Image Level *i*DataAgent to restore directly to a virtual server's volumes.
- For the Image Level *i*DataAgent on Unix, Unix File System snapshots cannot be backed up by non-CXBF backup jobs.
- You can not mix CXBF and Checksum backups.You can not perform a Checksum incremental followed by a CXBF full backup or vice versa. Once you configure a volume as CXBF device, you cannot include it in the subclient configured for Checksum backups.

Back to Top

Backup - Image Level - How To

Topics | How To | Related Topics

RUN A CXBF BACKUP

Related Topics

- Subclients - SAN iDataAgents

Required Capability: See Capabilities and Permitted Actions

To run a CXBF backup:

1. Ensure that QSnap is installed on the client.
2. Configure the subclient.
3. From the Subclient Properties (General) dialog box, select the **CXBF** option.
4. Run the backup. (Select from any of the appropriate procedures below to start or schedule a backup for instructions.)

RUN A NON-CXBF BACKUP

Related Topics

- Subclients - SAN iDataAgents
- Pre/Post Processes

Required Capability: See Capabilities and Permitted Actions

To run a non-CXBF backup:

1. Configure the subclient.
2. From the Subclient Properties (General) dialog box, select the **Checksum** option.
3. Provide the appropriate PreScan and PostBackup scripts for snap and mount operations. For more information, see Pre/Post Processes for Data Protection Operations: Image Level and Image Level ProxyHost *i*DataAgents.
4. Run the backup. (Select from any of the appropriate procedures below to start or schedule a backup for instructions.)

BACK UP ORACLE INSTANCES WITH THE IMAGE LEVEL *I*DATAAGENT

You can back up Oracle instances with the Image Level *i*DataAgent. For more information, see Oracle BLI.

Required Capability: See Capabilities and Permitted Actions

To back up an Oracle Instance:

1. Create `PreOraScan` and `PostOraScan` scripts. See the Agent-Specific Guidelines for more information about these Pre/Post processing commands.
2. Create a Subclient which contains both the Oracle data files volume and the Oracle archive logs volume.
3. Specify the `PreOraScan` and `PostOraScan` scripts (created in Step 1) for the PreScan and PostScan phase. See Configure a Subclient for Pre/Post Processing.

   NOTES

   - It is recommended that you select the **Run the Post Scan Process for all the attempts** checkbox; in case of a snap failure, the Oracle DB would remain is suspended state until user intervention if this checkbox is not selected.
4. Enter the local user account name of the user who installed or created the Oracle instance on the machine. See Change Account for Executing Pre/Post Commands (Data Protection) for instructions.
5. Run the Image Level backup. (Select from any of the backup procedures below for instructions.)
6. Verify that the PreScan and PostScan phases succeeded. The `CVVIOrascan.log` will be created in the log file directory, which contains logs of snapping the Oracle archive log volume.

To verify the Oracle database is really quiesced:

In the `admin bdump` directory of the Oracle instance, open the `alert<instancename>.log` to verify that the Oracle instance was properly quiesced and unquiesced. This log file records all the activities of the Oracle instance.

START A BACKUP

Before You Begin

- Depending on your agent, you can perform the following types of backup operations: **Full**, **Incremental**, **Differential** or **Synthetic Full.**
  - Read Full Backups before performing a Full Backup.
  - Read Incremental Backups before performing a Incremental Backup.
  - Read Differential Backups before performing a Differential Backup.
  - Read Synthetic Full Backups before performing a Synthetic Full Backup.

Required Capability: See Capabilities and Permitted Actions

To start an immediate backup job:

1. From the CommCell Browser, expand **Client Computers** by double-clicking **Client Computers** | ***i*DA File System** | **defaultBackupSet**. The default and other subclients (if available) are displayed on the right-hand windowpane.
2. To back up the default subclient, right-click the subclient, and click **Backup**.
3. From the Backup Options dialog box, select **Run Immediately**.
4. Select **Full**, **Incremental**, **Differential** or **Synthetic Full** backup.

   In certain circumstances a non-full backup may automatically be converted to a full backup. For a listing of these circumstances, see When a Non-Full Backup is Automatically Converted to a Full Backup.
5. Click **OK**. You can track the progress of the backup job from the **Job Controller** window. If you are using a stand-alone drive, you are prompted to load a specific cartridge into the drive. If you are using a library, you will not receive this prompt. The system loads the tapes automatically. Your cartridges should be appropriately labeled. This will enable you to locate the correct cartridge for a restore job, if necessary.
6. When the backup has completed, the Job Controller displays `Completed`.

   After running a backup, you may want to verify the backup data. You can do this by viewing the Backup History. For more information, see Backup Job History.

   - You can also run backups of the following:
     - For a user-defined backup set or instance, right-click the backup set you want to back up, click **All Tasks**, and click **Backup All Subclients**.
     - For the Lotus Notes Document *i*DataAgent, to back up a partition, right-click the partition you want to back up, click **All Tasks**, and click **Backup Default Backup Set**.
       For the Lotus Notes Database *i*DataAgent, to back up a partition, right-click the partition you want to back up, click **All Tasks**, and click **Backup All Subclients**.
     - For Agents that do not have backup set or instance levels, to back up all subclients, right-click the agent icon, click **All Tasks**, and click **Backup All Subclients**.

- If you chose a level higher than subclient (i.e., backup set, etc.), you are prompted to confirm that you want to back up all the subclients below that level/node. Click **Yes**.
- Starting a data protection operation on a backup set, instance or agent level causes the system to start individual data protection operations for each subclient contained therein. If the subclients are associated with the same storage policy, then their operations will run sequentially unless that storage policy is configured to accommodate multiple data streams.

START A SYNTHETIC FULL BACKUP

Before You Begin

- Read Synthetic Full Backups before performing a Synthetic Full Backup.
- For SharePoint Document, for a versioned document that has multiple versions, all of the backed up versions can be viewed in the **View All Versions** window and restored, until a Synthetic Full backup is run. After running the Synthetic Full backup you can only view and restore the latest backed up version for the document.

Required Capability: See Capabilities and Permitted Actions

To start an immediate backup job:

1. From the CommCell Browser, expand **Client Computers** by double-clicking **Client Computers** | ***i*DA File System** | **defaultBackupSet**. The default and other subclients (if available) are displayed on the right-hand windowpane.
2. To back up the default subclient, right-click the subclient, and click **Backup**.
3. From the Backup Options dialog box, select **Run Immediately**.
4. Select **Synthetic Full** backup.

   Running an incremental backup immediately before the synthetic full ensures that any new or recently changed data is included in the synthetic full. Running an incremental backup immediately after the synthetic full ensures that any new or recently changed data since the backup that occurred prior to the synthetic full, but was not included in the synthetic full, is backed up by the incremental. Remember, a synthetic full consolidates data; it does not actually back up data from the client computer.
5. Click **OK**. You can track the progress of the backup job from the **Job Controller** window. If you are using a stand-alone drive, you are prompted to load a specific cartridge into the drive. If you are using a library, you will not receive this prompt. The system loads the tapes automatically. Your cartridges should be appropriately labeled. This will enable you to locate the correct cartridge for a restore job, if necessary.
6. When the backup has completed, the **Job Controller** displays `Completed`.

   After running a backup, you may want to verify the backup data. You can do this by viewing the Backup History. For more information, see Backup Job History.

You can also run synthetic full backups of the following:

- For a user-defined backup set or instance, right-click the backup set you want to back up, click **All Tasks**, and click **Backup All Subclients**.
- For the Lotus Notes Document *i*DataAgent, to back up a partition, right-click the partition you want to back up, click **All Tasks**, and click **Backup Default Backup Set**.
  For the Lotus Notes Database *i*DataAgent, to back up a partition, right-click the partition you want to back up, click **All Tasks**, and click **Backup All Subclients**.
- For Agents that do not have backup set or instance levels, to back up all subclients, right-click the agent icon, click **All Tasks**, and click **Backup All Subclients**.
  - If you chose a level higher than subclient (i.e., backup set, etc.), you are prompted to confirm that you want to back up all the subclients below that level/node. Click **Yes**.
  - Starting a data protection operation on a backup set, instance or agent level causes the system to start individual data protection operations for each subclient contained therein. If the subclients are associated with the same storage policy, then their operations will run sequentially unless that storage policy is configured to accommodate multiple data streams.

SCHEDULE BACKUPS

You can schedule backups to occur with the following procedure. You will be prompted to create a schedule for the data protection operation after selecting your data protection options.

Before You Begin

- **All Agents**
  - Be sure all of the subclients are backed up, or scheduled to be backed up as needed, in order to secure all of the data for the agent. Note this does not apply to archive operations.

Required Capability: See Capabilities and Permitted Actions

To schedule a backup operation:

1. From the CommCell Browser, select one of the following:
   - To back up a subclient, right-click the subclient and click **Backup**.
   - To back up a backup set or instance, right-click the backup set or instance, click **All Tasks**, and click **Backup All Subclients**.
   - To back up the default backup set, right-click the agent or instance node, click **All Tasks**, and click **Backup Default Backup Set**.
   - For the Lotus Notes Document *i*DataAgent, to back up a partition, right-click the partition you want to back up, click **All Tasks**, and click **Backup Default Backup Set**.
     For the Lotus Notes Database *i*DataAgent, to back up a partition, right-click the partition you want to back up, click **All Tasks**, and click **Backup All Subclients**.
   - For Agents that do not have backup set or instance levels, to back up all subclients, right-click the agent icon, click **All Tasks**, and click **Backup All Subclients**.
2. If you chose a level higher than subclient (i.e., backup set, etc.), you are prompted to confirm that you want to back up all the subclients below that level/node. Click **Yes**.
3. From the Backup Options dialog box, select the type of backup that you want to initiate. In certain circumstances a non-full backup may automatically be converted to a full backup. For a listing of these circumstances, see When a Non-Full Backup is Automatically Converted to a Full Backup.
4. Click **Schedule**. Click **OK** to continue.
5. From the Schedule Details (Schedule Details) dialog box, create a schedule for this operation. For step-by-step instructions, see Create a Job Schedule. Click **OK** to continue.
6. Your backup operation will execute according to the specified schedule.

   Starting a data protection operation on a backup set, instance or agent level causes the system to start individual data protection operations for each subclient contained therein. If the subclients are associated with the same storage policy, then their operations will run sequentially unless that storage policy is configured to accommodate multiple data streams.

START A BACKUP IN THE SUSPENDED STATE

Use the following procedure to start a backup in the suspended state.

Before You Begin

- Be sure all of the subclients are backed up, or scheduled to be backed up as needed, in order to secure all of the data for the agent. Note this does not apply to archive operations.

Required Capability: See Capabilities and Permitted Actions

To start an immediate backup job with advanced backup options:

1. From the CommCell Browser, select one of the following:
   - To backup a subclient, right-click the subclient to want to backup and click **Backup**.
   - To backup a user-defined backup set or instance, right-click the backup set you want to backup, click **All Tasks**, and click **Backup All Subclients**.
   - To backup the default backup set, right-click the agent or instance node, click **All Tasks**, and click **Backup Default Backup Set**.
   - For Lotus Notes *i*DataAgent, to backup a partition, right-click the partition you want to backup, click **All Tasks**, and click **Backup Default Backup Set**.
   - For Agents that do not have backup set or instance levels, to backup all subclients, right-click the agent icon, click **All Tasks**, and click **Backup All Subclients**.
2. If you chose a level higher than subclient (i.e., backup set, etc.), you are prompted to confirm that you want to back up all the subclients below that level/node. Click **Yes**.
3. From the Backup Options dialog box, select **Run Immediately**.
4. Select the type of backup that you want to initiate.
   In certain circumstances a non-full backup may automatically be converted to a full backup. For a listing of these circumstances, see When a Non-Full Backup is Automatically Converted to a Full Backup.
5. Click the **Advanced** button to open the **Advanced Backup Options** dialog box.
6. Click on the Advanced Backup Options (Startup) tab, and select **Startup in suspended state** and click **OK**.
7. From the **Backup Options** dialog box, click **OK**. You can track the progress of the backup job from the Job Controller window.
8. If you are using a stand-alone drive, you are prompted to load a specific cartridge into the drive. If you are using a library, you will not receive this prompt. The system loads the tapes automatically.

   Your cartridges should be appropriately labeled. This will enable you to locate the correct cartridge for a restore job,

if necessary.

9. When the backup has completed, Job Controller displays `Completed`.

   Starting a data protection operation on a backup set, instance or agent level causes the system to start individual data protection operations for each subclient contained therein. If the subclients are associated with the same storage policy, then their operations will run sequentially unless that storage policy is configured to accommodate multiple data streams.

START A BACKUP ON NEW MEDIA

Use the following procedure to start a backup on new media.

Before You Begin

- Be sure all of the subclients are backed up, or scheduled to be backed up as needed, in order to secure all of the data for the agent. Note this does not apply to archive operations.

Required Capability: See Capabilities and Permitted Actions

To start an immediate backup job with advanced backup options:

1. From the CommCell Browser, select one of the following:
   - To backup a subclient, right-click the subclient to want to backup and click **Backup**.
   - To backup a user-defined backup set or instance, right-click the backup set you want to backup, click **All Tasks**, and click **Backup All Subclients**.
   - To backup the default backup set, right-click the agent or instance node, click **All Tasks**, and click **Backup Default Backup Set**.
   - For Lotus Notes *i*DataAgent, to backup a partition, right-click the partition you want to backup, click **All Tasks**, and click **Backup Default Backup Set**.
   - For Agents that do not have backup set or instance levels, to backup all subclients, right-click the agent icon, click **All Tasks**, and click **Backup All Subclients**.
2. If you chose a level higher than subclient (i.e., backup set, etc.), you are prompted to confirm that you want to back up all the subclients below that level/node. Click **Yes**.
3. From the Backup Options dialog box, select **Run Immediately**.
4. Select the type of backup that you want to initiate.
   In certain circumstances a non-full backup may automatically be converted to a full backup. For a listing of these circumstances, see When a Non-Full Backup is Automatically Converted to a Full Backup.
5. Click the **Advanced** button to open the **Advanced Backup Options** dialog box.
6. Click on the Advanced Backup Options (Media) tab, and select **Start new media** and click **OK**.
   If you would like jobs with other Job IDs to use this new media, also select the **Allow other schedule to use media set** option.
7. From the **Backup Options** dialog box, click **OK**. You can track the progress of the backup job from the Job Controller window.
8. If you are using a stand-alone drive, you are prompted to load a specific cartridge into the drive. If you are using a library, you will not receive this prompt. The system loads the tapes automatically.

   Your cartridges should be appropriately labeled. This will enable you to locate the correct cartridge for a restore job, if necessary.

9. When the backup has completed, Job Controller displays `Completed`.

   Starting a data protection operation on a backup set, instance or agent level causes the system to start individual data protection operations for each subclient contained therein. If the subclients are associated with the same storage policy, then their operations will run sequentially unless that storage policy is configured to accommodate multiple data streams.

START A BACKUP THAT MARKS MEDIA FULL ON COMPLETION

Use the following procedure to start a backup that marks media full on completion.

Before You Begin

- Be sure all of the subclients are backed up, or scheduled to be backed up as needed, in order to secure all of the data for the agent. Note this does not apply to archive operations.

Required Capability: See Capabilities and Permitted Actions

To start an immediate backup job with advanced backup options:

1. From the CommCell Browser, select one of the following:
   - To backup a subclient, right-click the subclient to want to backup and click **Backup**.
   - To backup a user-defined backup set or instance, right-click the backup set you want to backup, click **All Tasks**, and click **Backup All Subclients**.
   - To backup the default backup set, right-click the agent or instance node, click **All Tasks**, and click **Backup Default Backup Set**.
   - For Lotus Notes *i*DataAgent, to backup a partition, right-click the partition you want to backup, click **All Tasks**, and click **Backup Default Backup Set**.
   - For Agents that do not have backup set or instance levels, to backup all subclients, right-click the agent icon, click **All Tasks**, and click **Backup All Subclients**.
2. If you chose a level higher than subclient (i.e., backup set, etc.), you are prompted to confirm that you want to back up all the subclients below that level/node. Click **Yes**.
3. From the Backup Options dialog box, select **Run Immediately**.
4. Select the type of backup that you want to initiate.
   In certain circumstances a non-full backup may automatically be converted to a full backup. For a listing of these circumstances, see When a Non-Full Backup is Automatically Converted to a Full Backup.
5. Click the **Advanced** button to open the **Advanced Backup Options** dialog box.
6. Click on the Advanced Backup Options (Media) tab, and select **Mark media full after successful operation** and click **OK**.
7. From the **Backup Options** dialog box, click **OK**. You can track the progress of the backup job from the Job Controller window.
8. If you are using a stand-alone drive, you are prompted to load a specific cartridge into the drive. If you are using a library, you will not receive this prompt. The system loads the tapes automatically.

   Your cartridges should be appropriately labeled. This will enable you to locate the correct cartridge for a restore job, if necessary.

9. When the backup has completed, Job Controller displays `Completed`.

   Starting a data protection operation on a backup set, instance or agent level causes the system to start individual data protection operations for each subclient contained therein. If the subclients are associated with the same storage policy, then their operations will run sequentially unless that storage policy is configured to accommodate multiple data streams.

START A BACKUP THAT SKIPS METAFILE CREATION

Use the following procedure to start an Image Level or Image Level ProxyHost *i*DataAgent backup that skips metafile creation. For more information, see one of the following:

- Skip Metafile Creation - Image Level
- Skip Metafile Creation - Image Level ProxyHost

Before You Begin

- Be sure all of the subclients are backed up, or scheduled to be backed up as needed, in order to secure all of the data for the agent.
- If File Level restores will never be used for this client, instead of manually selecting the Skip Metafile Creation option for each job, edit the `SkipMetaFileCreation` registry key to automatically skip metafile creation for all backup jobs.

Required Capability: See Capabilities and Permitted Actions

To start an immediate backup job with advanced backup options:

1. Select one of the following:
   - To backup a subclient, right-click the subclient to want to backup and click **Backup**.
   - To backup a user-defined backup set, right-click the backup set you want to backup, click **All Tasks**, and click **Backup All Subclients**.
   - To backup the default backup set, right-click the agent or instance node, click **All Tasks**, and click **Backup Default Backup Set**.
2. If you chose a level higher than subclient (i.e., backup set, etc.), you are prompted to confirm that you want to back up all the subclients below that level/node. Click **Yes**.
3. From the Backup Options dialog box, select **Run Immediately**.
4. Select the type of backup that you want to initiate.
   In certain circumstances a non-full backup may automatically be converted to a full backup. For a listing of these circumstances, see When a Non-Full Backup is Automatically Converted to a Full Backup.
5. Click the **Advanced** button to open the **Advanced Backup Options** dialog box.

6. Click on the Advanced Backup Options (Media) tab, and select **Skip Metafile Creation** and click **OK**.
7. From the **Backup Options** dialog box, click **OK**. You can track the progress of the backup job from the Job Controller window.
8. If you are using a stand-alone drive, you are prompted to load a specific cartridge into the drive. If you are using a library, you will not receive this prompt. The system loads the tapes automatically.

   Your cartridges should be appropriately labeled. This will enable you to locate the correct cartridge for a restore job, if necessary.

9. When the backup has completed, Job Controller displays `Completed`.

   Starting a backup on a backup set, instance or agent level causes the system to start individual backup jobs for each subclient contained therein. If the subclients are associated with the same storage policy, then their jobs will run sequentially unless that storage policy is configured to accommodate multiple data streams.

START A BACKUP THAT RELEASES RESOURCES DURING THE META-DATA COLLECTION PHASE

Use the following procedure to start a backup that Releases Resources during the meta-data collection phase.

Before You Begin

- **All Agents**
  - Be sure all of the subclients are backed up, or scheduled to be backed up as needed, in order to secure all of the data for the agent.
- **Image Level**
  - If File Level restores will never be used for this client, instead of manually selecting the Skip Metafile Creation option for each job, edit the `SkipMetaFileCreation` registry key to automatically skip metafile creation for all backup jobs.

Required Capability: See Capabilities and Permitted Actions

To start an immediate backup job with advanced backup options:

1. Select one of the following:
   - To backup a subclient, right-click the subclient to want to backup and click **Backup**.
   - To backup a user-defined backup set, right-click the backup set you want to backup, click **All Tasks**, and click **Backup All Subclients**.
   - To backup the default backup set, right-click the agent or instance node, click **All Tasks**, and click **Backup Default Backup Set**.
2. If you chose a level higher than subclient (i.e., backup set, etc.), you are prompted to confirm that you want to back up all the subclients below that level/node. Click **Yes**.
3. From the Backup Options dialog box, select **Run Immediately**.
4. Select the type of backup that you want to initiate.
   In certain circumstances a non-full backup may automatically be converted to a full backup. For a listing of these circumstances, see When a Non-Full Backup is Automatically Converted to a Full Backup.
5. Click the **Advanced** button to open the **Advanced Backup Options** dialog box.
6. Click on the Advanced Backup Options (Media) tab, and select **Release Resources during meta-data collection phase** and click **OK**.
7. From the **Backup Options** dialog box, click **OK**. You can track the progress of the backup job from the Job Controller window.
8. If you are using a stand-alone drive, you are prompted to load a specific cartridge into the drive. If you are using a library, you will not receive this prompt. The system loads the tapes automatically.

   Your cartridges should be appropriately labeled. This will enable you to locate the correct cartridge for a restore job, if necessary.

9. When the backup has completed, Job Controller displays Completed.

   All Agents

   Starting a backup on a backup set, instance or agent level causes the system to start individual backup jobs for each subclient contained therein. If the subclients are associated with the same storage policy, then their jobs will run sequentially unless that storage policy is configured to accommodate multiple data streams.

START A BACKUP THAT RESERVES RESOURCES BEFORE SCAN

Use the following procedure to start a backup that reserves the backup media before the scan phase.

Before You Begin

- Be sure all of the subclients are backed up, or scheduled to be backed up as needed, in order to secure all of the data for the agent.

Required Capability: See Capabilities and Permitted Actions

To start an immediate backup job with advanced backup options:

1. select one of the following:
   - To backup a subclient, right-click the subclient to want to backup and click **Backup**.
   - To backup a user-defined backup set, right-click the backup set you want to backup, click **All Tasks**, and click **Backup All Subclients**.
   - To backup the default backup set, right-click the agent or instance node, click **All Tasks**, and click **Backup Default Backup Set**.
   - For Lotus Notes *i*DataAgent, to backup a partition, right-click the partition you want to backup, click **All Tasks**, and click **Backup Default Backup Set**.
   - For the Microsoft SQL Server *i*DataAgent, to backup a database, right-click the database you want to backup, click **All Tasks**, and click **Backup Database**.
   - For the Microsoft SQL Server *i*DataAgent, to backup up an instance, right-click the instance you want to backup, click **All Tasks**, and click **Backup SQL Server**.
   - For Agents that do not have backup set or instance levels, to backup all subclients, right-click the agent icon, click **All Tasks**, and click **Backup All Subclients**.
2. If you chose a level higher than subclient (i.e., backup set, etc.), you are prompted to confirm that you want to back up all the subclients below that level/node. Click **Yes**.
3. From the Backup Options dialog box, select **Run Immediately**.
4. Select the type of backup that you want to initiate.
   In certain circumstances a non-full backup may automatically be converted to a full backup. For a listing of these circumstances, see When a Non-Full Backup is Automatically Converted to a Full Backup.
5. Click the **Advanced** button to open the **Advanced Backup Options** dialog box.
6. Click on the Advanced Backup Options (Media) tab, and select **Reserve Resources before scan** and click **OK**.
7. From the **Backup Options** dialog box, click **OK**. You can track the progress of the backup job from the Job Controller window.
8. If you are using a stand-alone drive, you are prompted to load a specific cartridge into the drive. If you are using a library, you will not receive this prompt. The system loads the tapes automatically.

   Your cartridges should be appropriately labeled. This will enable you to locate the correct cartridge for a restore job, if necessary.

9. When the backup has completed, Job Controller displays `Completed`.

   Starting a backup on a backup set, instance or agent level causes the system to start individual backup jobs for each subclient contained therein. If the subclients are associated with the same storage policy, then their jobs will run sequentially unless that storage policy is configured to accommodate multiple data streams.

START A BACKUP WITH A SET JOB PRIORITY

This option allows you to manually set a job priority. This is useful if you have jobs that are very important and must complete, and/or jobs that can be moved to a lower priority. For more information, see Job Priorities and Priority Precedence.

Before You Begin

- Be sure all of the subclients are backed up, or scheduled to be backed up as needed, in order to secure all of the data for the agent. Note this does not apply to archive operations.

Required Capability: See Capabilities and Permitted Actions

To start an immediate backup job with advanced backup options:

1. From the CommCell Browser, select one of the following:
   - To backup a subclient, right-click the subclient to want to backup and click **Backup**.
   - To backup a user-defined backup set or instance, right-click the backup set you want to backup, click **All Tasks**, and click **Backup All Subclients**.
   - To backup the default backup set, right-click the agent or instance node, click **All Tasks**, and click **Backup Default Backup Set**.
   - For Lotus Notes *i*DataAgent, to backup a partition, right-click the partition you want to backup, click **All Tasks**, and click **Backup Default Backup Set**.
   - For Agents that do not have backup set or instance levels, to backup all subclients, right-click the agent icon, click **All Tasks**, and click **Backup All Subclients**.

2. If you chose a level higher than subclient (i.e., backup set, etc.), you are prompted to confirm that you want to back up all the subclients below that level/node. Click **Yes**.
3. From the Backup Options dialog box, select **Run Immediately**.
4. Select the type of backup that you want to initiate.
   In certain circumstances a non-full backup may automatically be converted to a full backup. For a listing of these circumstances, see When a Non-Full Backup is Automatically Converted to a Full Backup.
5. Click the **Advanced** button to open the **Advanced Backup Options** dialog box.
6. Click on the Advanced Backup Options (Startup) tab, and select **Change Priority** and then enter a value. Click **OK** to continue.
7. From the **Backup Options** dialog box, click **OK**. You can track the progress of the backup job from the Job Controller window.
8. If you are using a stand-alone drive, you are prompted to load a specific cartridge into the drive. If you are using a library, you will not receive this prompt. The system loads the tapes automatically.

   Your cartridges should be appropriately labeled. This will enable you to locate the correct cartridge for a restore job, if necessary.

9. When the backup has completed, Job Controller displays `Completed`.

   Starting a data protection operation on a backup set, instance or agent level causes the system to start individual data protection operations for each subclient contained therein. If the subclients are associated with the same storage policy, then their operations will run sequentially unless that storage policy is configured to accommodate multiple data streams.

START A BACKUP WITH VAULT TRACKING ENABLED

Use the following procedure to start a backup with Vault Tracking enabled.

For additional information, see the following:

- VaultTracker
- VaultTracker Enterprise

Before You Begin

- Be sure all of the subclients are backed up, or scheduled to be backed up as needed, in order to secure all of the data for the agent.

Required Capability: See Capabilities and Permitted Actions

To start an immediate backup job with advanced backup options:

1. select one of the following:
   - To backup a subclient, right-click the subclient to want to backup and click **Backup**.
   - To backup a user-defined backup set or instance, right-click the backup set you want to backup, click **All Tasks**, and click **Backup All Subclients**.
   - To backup the default backup set, right-click the agent or instance node, click **All Tasks**, and click **Backup Default Backup Set**.
   - For Lotus Notes *i*DataAgent, to backup a partition, right-click the partition you want to backup, click **All Tasks**, and click **Backup Default Backup Set**.
   - For Agents that do not have backup set or instance levels, to backup all subclients, right-click the agent icon, click **All Tasks**, and click **Backup All Subclients**.
2. If you chose a level higher than subclient (i.e., backup set, etc.), you are prompted to confirm that you want to back up all the subclients below that level/node. Click **Yes**.
3. From the Backup Options dialog box, select **Run Immediately**.
4. Select the type of backup that you want to initiate.
   In certain circumstances a non-full backup may automatically be converted to a full backup. For a listing of these circumstances, see When a Non-Full Backup is Automatically Converted to a Full Backup.
5. Click the **Advanced** button to open the **Advanced Backup Options** dialog box.
6. Click on the **Vault Tracking** tab, and select the vault tracking options you want to use and click **OK**.
7. From the **Backup Options** dialog box, click **OK**. You can track the progress of the backup job from the Job Controller window.
8. If you are using a stand-alone drive, you are prompted to load a specific cartridge into the drive. If you are using a library, you will not receive this prompt. The system loads the tapes automatically.

   Your cartridges should be appropriately labeled. This will enable you to locate the correct cartridge for a restore job, if necessary.

9. When the backup has completed, Job Controller displays `Completed`.

Back To Top

Restore Data - Image Level

Topics | How To | Related Topics

OVERVIEW

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

The Image Level *i*DataAgent performs the following types of restore operations:

- File Level Restore - You can browse the files/folders in your backup and select the files/folders you want to restore.
- Volume Level Restore - You can browse the volumes you have backed up and select a volume to restore. The following restore options for volumes are available for the Image Level *i*DataAgent. These options are available from the Restore Options dialog box.
  - Physical Volumes - This option enables you to restore the selected content as a physical volume.
  - VMDK files (Windows only) - This option enables you to restore the selected volume as a virtual machine file.
  - Virtual Hard Disk file (Windows only) - This option enables you to restore the selected content as a virtual hard disk file.

When restoring data, you can, if desired, restore the data to a file system type that differs from the type in which it originated. For example, you can restore NTFS data to a FAT file system and restore FAT data to an NTFS file system. FAT file systems do not support Discretionary Access Control Lists (DACL); therefore, any NTFS data that you restore to a FAT partition loses its original access privileges. Conversely, when FAT file system data is restored to an NTFS file system, the restored data inherits the DACL of the destination directory.

Back to Top

RESTORE CONSIDERATIONS FOR THIS AGENT

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

GENERAL

- Review the general restore requirements described in What You Need to Know Before Performing a Restore.
- Since Image Level restores a portion of Volume Information to the MediaAgent Index Cache, if encryption was enabled when the data was backed up, the Volume Information will be encrypted as well. This requires a Pass-Phrase on the MediaAgent to decrypt during the restore. There are two ways to handle this:
  - Export an Encryption Pass-Phrase to the MediaAgent in use as well as the destination Client (if using Alternate Data Paths (GridStor), it needs to be exported to all the MediaAgents.)
  - Manually type the Pass-Phrase for each restore. See Recovering Encrypted Data (With a Pass-Phrase).
- For increased logging of activities during data recovery operations, the `dEnableIRestoreLog` registry key can be created.
- Although you can use the Image Level iDataAgent to back up volumes of a clustered shared disk, you cannot use the Image Level iDataAgent to restore directly to the volumes of a clustered shared disk.

FILE LEVEL RESTORE:

- You cannot restore any archived files and folders.
- It is recommended that file-level restores from disk or volume-level backups be performed only with small files. For example, restoring a 2GB file from a disk-level backup is not recommended.
- For a File Level Restore, you must select the **File Level Browse** option in the **Browse Options** dialog box.
- The Image Level *i*DataAgent on HP-UX does not support file level restores.
- Supported only for certain operating systems and file system types; see Supported Data Types.

- The Image Level and Image Level ProxyHost *i*DataAgents on Unix can perform a File Level Restore only when the OS of the Client computer and the MediaAgent are the same.
- For a Windows MediaAgent, the Index Cache folder must reside on an NTFS partition.
- If metadata was not collected during a backup, a File Level restore cannot be performed. This might result from having selected the Skip Metafile Creation option in the Advanced Backup Options screen, or as a result of failure to collect metadata during the backup, in which case the system generated an Event Message warning of the failure.
- Do not restore:
  - hidden objects
  - system files
  - recycler files
- Image Browse is not supported. A no-image browse operation returns the most recent version of the data that existed back to the most recent full backup, rather than returning an image of the specified entity (i.e., file system/directory). To perform a File Level Restore of an entity to a point in time, you will need to Browse back in time, to the point in time you want to restore.
- When a File Level restore is started, metadata is restored prior to the actual restore of files, and during this time, a message might indicate the Index Cache is being restored, although it already exists. Also, the media containing this metadata will be mounted into the library even if this media does not contain the data from the full backup.
- As noted in the Image Level Backup Considerations, if the cluster size (allocation unit) on the disk that you backed up was less than 1024 bytes, a file level restore job will not complete successfully.

VOLUME LEVEL RESTORE:

- For a Volume Level Restore, you must select the **Volume Level Browse** option in the **Browse Options** dialog box.
- Do not run Volume Level restore if destination volume contains OS files.
- For a Volume Level Restore from a Unix Checksum backup, volume detection must be run from Volume Explorer prior to the restore. The **Restore Options** dialog box will display the list of volumes available for the restore.
- When performing a Volume Level Restore, the destination volume will adopt the file system type of the restored data. Thus, for instance, if you restore an NTFS volume to a FAT destination volume, the destination volume will be NTFS as a result of the Volume Level Restore.
- When restoring a volume changes its file system type (see Supported Data Types), sometimes Windows Explorer may still show the volume having the same file system type as before the restore, even though Computer Management shows the new (correct) type. After a reboot, Windows Explorer will show the correct file system type.
- The destination volume must be at least as large as the volume from which the data was backed up. This is true regardless of the amount of data that is actually restored. For example, if you back up a 10 GB volume that contains 100 MB of data, you can only restore that data to a volume that is 10 GB or larger. The restore operation will fail if the destination volume is smaller than the source volume. We recommend that you restore data to a volume that is at least 1 MB larger than the source volume.

  Windows Logical Disk Manager (LDM) displays the size of a volume in round numbers. Consequently, LDM may display the same size for two volumes with slightly different block counts. To get the exact size of a volume, open the Windows Explorer, right-click the volume, and select **Properties**. The volume's capacity in bytes is its exact size.

  If no destination volume is specified, the system attempts to restore to a volume on the **Destination Computer** with the same name as the one being restored. If no such volume exists, the restore operation fails.

Back to Top

RESTORE DESTINATIONS

By default, the Image Level *i*DataAgent restores a volume or file(s) to the client from which it originated; this is referred to as an in-place restore. If desired, you can also restore the data to a different Image Level client. Keep in mind the following considerations when performing such restores:

- The destination client must reside in the same CommCell as the client whose data was backed up.
- Each of these restore destination types are available for both Volume Level restores as well as File Level restores, both for Windows and Unix.

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

IN-PLACE RESTORE

- Same path/destination

OUT-OF-PLACE RESTORE

- Same path/destination
- Different path/destination

CROSS-PLATFORM RESTORES

- Same Operating System - Different Version

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

  - Windows 2000 data can be restored to Windows 2003
  - Windows 2003 data can be restored to Windows 2000

RESTORE TO NETWORK DRIVE/NFS-MOUNTED FILE SYSTEM

For File Level Restores only, besides restoring data to a client computer's local drive, you can also restore data to a UNC path (Windows) or an NFS-Mounted File System (Unix). (See Restore to Network Drive/NFS-Mounted File System for comprehensive information.)

PERFORMING A FILE LEVEL RESTORE

If the Image Level backup contains the Metadata information, you can perform a File Level restore operation on the Image Level Backup. By default, files are restored to the same folder. However, while restoring the files related to system state, it is recommended to perform an Out-of-Place restore. This will ensure that the existing files are not overwritten.

Use the following steps to perform a file level restore on an Image Level backup:

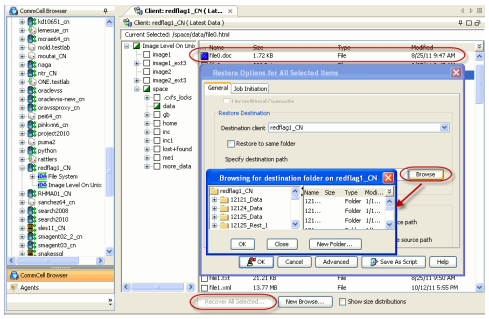

1. From the CommCell Browser, navigate to **Client Computers | *<Client>***.
2. Right-click ***<Image Level>***, point to **All Tasks**, and then click **Browse Backup Data**.
3. Click **OK**.
4. Navigate to the folder you want to restore and then select the files to be restored.
5. Click **Recover All Selected**.
6. Clear the **Restore to same folder** checkbox.
7. Specify the destination path by clicking the **Browse** button.

   This will ensure that the existing files are not overwritten.
8. Click **OK**.
9. Click the **Job Initiation** tab.
10. Click **OK**.

CHANGING THE TEMPORARY DIRECTORY FOR FILE LEVEL RESTORE

By default, the extents (blocks) are restored to the default temporary location, for example: `C:\Program Files\Company\Product\iDataAgent\FileSystemAgent`.

Based on the available size of recovery space on the target client, you can move the temporary directory to a new location using the following steps:

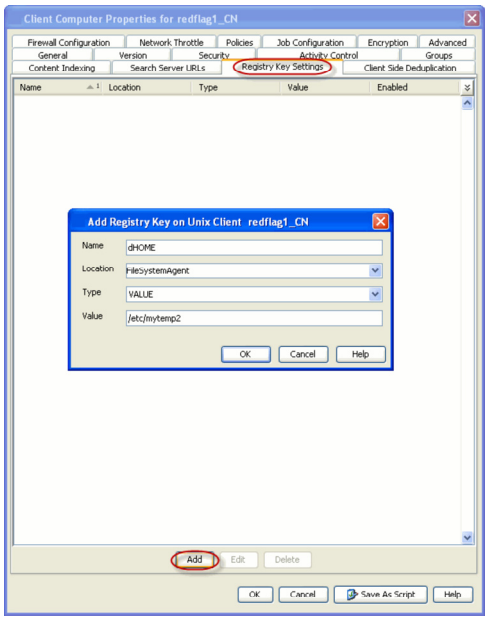

1. From the CommCell Console, navigate to **Client Computers**.
2. Right-click the ***<Client>***, and then click **Properties**.
3. Click the **Registry Key Settings** tab.
4. Click **Add**.
5. In the **Name** field, type `dHOME`.
6. In the **Location** list, type `FileSystemAgent`.
7. In the **Type** box:
   - on Unix client

     Select `VALUE`.
   - On Windows client

     Select `REG_SZ`.
8. In the **Value** field, type the path to the new directory.

   For example (on UNIX): `/etc/mytemp2`

   For example (on Windows): `D:\mytemp2`
9. Click **OK**.
10. Click **OK**.

RESTORE FROM BACKUP COPIES

The Image Level on Unix *i*DataAgent restores file(s) to the client from which it originated by default; this is referred to as an in-place restore. If desired, you can also restore the data to a different Image Level client.

Keep in mind the following considerations when performing such restores:

- The destination client must reside in the same CommCell as the client whose data was backed up.
- The Image Level on Unix iDataAgent can perform a file level restore only when the operating systems of the client computer and the MediaAgent are the same.

FILE LEVEL RESTORE FROM A PRIMARY BACKUP COPY

1. From the CommCell Console, navigate to **Client Computers | <*Client*>**.
2. Right-click the **Image Level on Unix** *i*DataAgent, and then click **All Tasks | Browse Backup Data**.
3. From the **Browse Options** dialog box, select the MediaAgent from the **Use MediaAgent** list, on which the primary copy of data is backed up.
4. Click **OK**.
5. Select the file/s you want to restore, and then click **Recover All Selected**.
6. In the **Destination Client** list, select the client with the same operating system as the MediaAgent.
7. Click **OK**.

FILE LEVEL RESTORE FROM AN AUXILIARY BACKUP COPY

Please note that if the operating systems of the client computer and the MediaAgent to which the primary backup copy is directed are different then you can perform a file level restore using the auxiliary backup copy.

1. From the CommCell Console, navigate to **Client Computers | <*Client*>**.
2. Right-click the **Image Level on Unix** *i*DataAgent, and then click **All Tasks | Browse Backup Data**.
3. From the **Browse Options** dialog box, select the MediaAgent from the **Use MediaAgent** list, on which the primary copy of data is backed up.
4. Click **Advanced**.
5. Select the **Browse from copy precedence** check box, and then in the **Copy Precedence** box, type or select 2 as the copy precedence number.
6. Click **OK** to close the **Advanced Browse Options** dialog box.
7. Click **OK**.
8. Select the file/s you want to restore, and then click **Recover All Selected**.
9. In the **Destination Client** list, select the client with the same operating system as the MediaAgent.
10. Click **OK**.

FREQUENTLY ASKED QUESTIONS

WHAT HAPPENS DURING A FILE LEVEL RESTORE?

File Level restores are carried out as extractions of content from the Image Level backup. The following sequence of events takes place during a File Level Restore:

1. During Browse, if the index is not available in the cache, a request is sent to the Media Agent to index the objects selected for restore.
2. The Master File Table of the backup is scanned to determine the number of blocks (extents) that contain the files to be restored.
3. The MediaAgent restores the selected files to the designated client and places them in the temporary directory.
4. The Restore process on the client extracts the file data from the restored extents and then restores file(s) from them.
5. Once the restore operation is complete, it will clean up the temporary directory to which the extents were restored.

> While restoring large number of files, you may not see the data being restored instantly as the files need to be extracted from the disk image.

HOW TO ESTIMATE THE FREE SPACE REQUIRED FOR EXTENT RECOVERY LOCATION?

The temporary space required for the extent recovery location would be the size of the files being restored with a small amount of overhead (10% - 20%). If

the disk was highly fragmented at the time of backup, this may require additional temporary space. See examples below:

Example 1: If the default Extent size is 1 MB, in order to restore a 4 KB file, a 1 MB extent is restored on the client to allow the extraction of the 4 KB file.

Example 2: When the disk is not highly fragmented at the time of backup, if you try to restore files that are collectively 1 GB in size, then 1.2 GB temporary space is required to hold the extents.

Example 3: If a 400 KB file is highly fragmented and scattered in 100 different extents, then all the 100 extents (100 MB) will need to be restored to the machine to restore the 400 KB file.

HOW DO THE DEFAULT CLUSTER SIZES FOR NTFS AFFECT RESTORES?

NTFS volumes allocate hard disk space using increments of cluster sizes. A cluster is a smallest fixed unit of disk space that can be allocated to a file. For file sizes that are not an exact multiple of the cluster size, additional space must be allocated as the next largest multiple of the cluster size.

If the cluster size is not specified when formatting a partition, defaults are used according to the size of the partition, to reduce the amount of unused space and reduce fragmentation. You can override the default settings when formatting a partition.

For example, the default maximum cluster size for NTFS under Windows NT 4.0 and later is four kilobytes and NTFS file compression is not supported on drives with a larger cluster size.

When restoring volumes, the backed up cluster size determines the maximum cluster size for the destination volume. For larger volumes, use a larger cluster size.

The following table shows the default cluster sizes for NTFS.

	Default cluster sizes for NTFS		
Volume size	Windows NT 3.51	Windows NT 4.0	Windows 7, Windows Server 2008 R2, Windows Server 2008, Windows Vista, Windows Server 2003, Windows XP, Windows 2000
7 MB–512 MB	512 bytes	4 KB	4 KB
512 MB–1 GB	1 KB	4 KB	4 KB
1 GB–2 GB	2 KB	4 KB	4 KB
2 GB–2 TB	4 KB	4 KB	4 KB
2 TB–16 TB	Not Supported*	Not Supported*	4 KB
16TB–32 TB	Not Supported*	Not Supported*	8 KB
32TB–64 TB	Not Supported*	Not Supported*	16 KB
64TB–128 TB	Not Supported*	Not Supported*	32 KB
128TB–256 TB	Not Supported*	Not Supported*	64 KB
> 256 TB	Not Supported	Not Supported	Not Supported

* Not supported because of the limitations of the master boot record (MBR).

Windows limits the size of an NTFS volume to that addressable with 32-bit clusters, which is slightly less than 256 TB (using 64-KB clusters).

For more information, refer to Microsoft KB article 140365.

FILE LEVEL RESTORES AND METADATA COLLECTION ON AIX

Consider the following when doing a file level restore:

- If the volume has a large number of small files (>5 million) to be restored, then turn off metadata collection, do a volume level restore, and copy the desired files back instead of running a file level restore job.
- If the volume has a mixture of large and small files then use Scan Optimization for metadata collection.
- If you are backing up multiple disks from the same subclient then use multi-streaming to improve performance.
- If there will never be a file level restore then it is recommended to turn off metadata collection.

Performance Details:

OS version	AIX 6.1
dSnapChunkSize	65536
LUN Size	1.5 TB
Memory	8 GB
Processors	2

Metadata Collection Performance:

Total Number of Files : 107862987

JOB ID	**METADATA COLLECTION**	**SCAN TIME**	**BACKUP TIME**	**METADATA COLLECTION**	**ARCHIVE INDEX TIME**	**CLEANUP TIME**	**BACKUP SIZE**

	METHOD			TIME			
48	Using Scan Optimization	00:03:14	03:50:17	04:54:17	00:01:43	00:00:03	516.92 MB
56	Without Scan Optimization	00:00:24	03:42:58	18:54:20	00:01:54	00:00:03	516.92 MB
64	No Metadata Collection	00:00:21	03:50:48	00:00:07	00:00:15	00:00:02	516.92 MB

Backup Performance:

- With Scan Optimization Enabled

JID	JOB TYPE (F OR I)	NUMBER OF STREAMS	APPROXIMATE SIZE OF DATA	NUMBER OF FILES	SIZE OF APP.	DATA WRITTEN	NUMBER OF EXTENTS	SCAN TIME	BACKUP TIME	METADATA COLLECTION TIME	THROUGHP
834	F	4	1.02 TB+	7555148	1,106.71 GB	1,110.11 GB	566, 647	00:44:42	24:51:00	00.24:22	44.53 GB/h
843	I	4	100 GB+	743032	129.98 GB	131.18 GB	66,559	00:02:39	02:40:25	00:32:34	48.90 GB/h
844	I	4	100 GB+	743032	117.84 GB	119.09 GB	60,341	00:02:03	02:39:42	00:38:38	44.10 GB/h

- Without Scan Optimization

JID	JOB TYPE (F OR I)	NUMBER OF STREAMS	APPROXIMATE SIZE OF DATA	NUMBER OF FILES	SIZE OF APP.	DATA WRITTEN	NUMBER OF EXTENTS	SCAN TIME	BACKUP TIME	METADATA COLLECTION TIME	THROUGHP
847	F	4	1.02 TB+	7555148	1,146.00 GB	1,149.50 GB	586,765	00:02:18	24:04:50	01.22:44	47.59 GB/h
849	I	4	100 GB+	743032	130.76 GB	131.97 GB	66,959	00:02:04	02:39:55	01:34:42	49.19 GB/h
850	I	4	100 GB+	743032	120.77 GB	122.02 GB	61,840	00:01:18	02:19:35	02:02:52	51.82 GB/h

- Single Stream and No Scan Optimization

JID	JOB TYPE (F OR I)	NUMBER OF STREAMS	APPROXIMATE SIZE OF DATA	NUMBER OF FILES	SIZE OF APP.	DATA WRITTEN	NUMBER OF EXTENTS	SCAN TIME	BACKUP TIME	METADATA COLLECTION TIME	THROUGHP
858	F	1	1.02 TB+	7555148	1,146.00 GB	1,149.50 GB	586,753	00:00:20	32:00:12	01.23:04	35.81 GB/h
859	I	1	100 GB+	743032	121.16 GB	122.27 GB	62,034	00:02:13	04:00:19	01:24:20	30.43 GB/h
860	I	1	100 GB+	743032	141.77 GB	144.13 GB	72,590	00:01:44	03:58:04	01:30:39	28.68 GB/h

Back to Top

Restore Data - Image Level - How To

Topics | How To | Related Topics

RESTORE AN ORACLE INSTANCE USING IMAGE LEVEL

Before you Begin

Review the general and agent-specific restore requirements accessed from Restore Backup Data prior to performing any restore operation.

Required Capability: See Capabilities and Permitted Actions

To restore an Oracle Instance:

1. Shut down the Oracle database:

   ```
   SQL> shutdown immediate
   ```

2. Perform a Browse and Restore procedure to restore all volumes of the Oracle Instance.
3. In the restored Oracle archive log directory, locate the following two files:
   - `CTRLFILELOCATIONS` ----> This file contains the location of the all the control files.
   - `CTRLFILEBACKUP.CTL` ----> This is the backup control file.
4. Use `CTRLFILEBACKUP.CTL` to overwrite all the control files mentioned in the `CTRLFILELOCATIONS`.

   Example for Windows; if your restored Oracle archive log volume is `g:`, then type the following commands:

   - `copy g:\oracleDB\CtrlFileBackup.ctl c:\oracledata\control01.ctl`
   - `copy g:\oracleDB\CtrlFileBackup.ctl c:\oracledata\control02.ctl`
   - `copy g:\oracleDB\CtrlFileBackup.ctl c:\oracledata\control03.ctl`
5. Log in to SQL with `sysdba` privileges, and type the following command:

   ```
   startup mount;
   ```

6. After the database is mounted, type the following command to recover the database:

   ```
   recover database until cancel using backup controlfile;
   ```

7. After the prompt, type `auto` or enter the archive log directory location to apply the archived logs.
8. Type the following command:

   ```
   alter database open resetlogs;
   ```

9. Now the database should be in open mode; the production server is now fully restored and ready for use.

RESTORE AN ORACLE INSTANCE OUT-OF-PLACE USING IMAGE LEVEL

Before you Begin

Review the general and agent-specific restore requirements accessed from Restore Backup Data prior to performing any restore operation.

Required Capability: See Capabilities and Permitted Actions

To restore an Oracle Instance out-of-place:

1. Perform a Browse and Restore procedure to restore all volumes of the Oracle Instance.
2. After the restore job has completed, shut down the Oracle database:
3. Change the drive letter or mount point of the restored Oracle data file volume and Oracle archive file volume, to match the original drive letter or mount point of the Oracle volumes.
4. In the restored Oracle archive log directory, locate the following two files:
   - `CTRLFILELOCATIONS` ----> This file contains the location of the all the control files.
   - `CTRLFILEBACKUP.CTL` ----> This is the backup control file.

5. Use `CTRLFILEBACKUP.CTL` to overwrite all the control files mentioned in `CTRLFILELOCATIONS`.

   Example for Windows; if your restored Oracle archive log volume is `g:`, then type the following commands:

   - `copy g:\oracleDB\CtrlFileBackup.ctl c:\oracledata\control01.ctl`
   - `copy g:\oracleDB\CtrlFileBackup.ctl c:\oracledata\control02.ctl`
   - `copy g:\oracleDB\CtrlFileBackup.ctl c:\oracledata\control03.ctl`

6. Log in to SQL with `sysdba` privileges, and type the following command:

   ```
   startup mount;
   ```

7. After the database is mounted, type the following command to recover the database:

   ```
   recover database until cancel using backup controlfile;
   ```

8. After the prompt, type the following:

   ```
   auto
   ```

9. Type the following command:

   ```
   alter database open resetlogs;
   ```

10. Now the database should be in open mode; the production server is now fully restored and ready for use.

Back To Top

Basic Restore

Topics | How To | Related Topics

OVERVIEW

There are two functions that help you retrieve backed up data from the backup media: Browse and Restore. Browse operations allow you to view the data that has been backed up for a client computer without actually restoring the data. Restore operations retrieve the data from the backup media and restore it to the desired location.

In the CommCell Browser, the Browse and various Restore commands appear in the right-click menus at the agent, instance and/or backup set levels, depending on agent.

Using the Restore commands, i.e., restoring without browsing, is most appropriate when you want to restore the latest backup job for an agent, instance or backup set and want to retain the current file structure.

In certain situations and for supported agents, Restore operations can run without utilizing the Browse feature. For example, if you know the path/name of the volume of the data that you want to restore, you can restore it without browsing. In these agents, this procedure is most appropriate when the number of paths for the data that you want to restore is small or when the data that you want to restore is at a single volume. If you want to restore data from many different paths or volumes, you should probably select the data from the Browse window.

TIME RANGE OPTIONS

In a Basic Restore, point-in-time data can be restored by specifying time range options in the Advanced Restore Options (Time Range) dialog box. See Time Range Options for Basic Restore for step-by-step instructions. When a time range is specified for a Basic Restore, the latest version of the backup job within the given time range will be restored.

Basic Restore - How To

Topics | How To | Related Topics

BASIC RESTORE

Before You Begin:

- Review the general and agent-specific restore requirements accessed from Restore Backup Data prior to performing any restore operation.

Required Capability: See Capabilities and Permitted Actions

To restore data without using browse:

1. From the CommCell Browser, right-click the agent, instance or backup set that contains the data you want to restore, click **All Tasks** and then click the available **Restore** command (command names vary by agent).
2. In the **Restore** dialog box, type the starting path of the data you want to restore.
3. From the **Restore Options** and **Advanced Restore Options** dialog boxes, select the restore options that you want to use. When you accept all the default settings, you will be restoring the selected data to its original location. See Restore Backup Data for complete information on the agent-specific restore options.
4. After completing your selections, you can either start an immediate restore or schedule the restore.
   - If you want to schedule the job, click the Job Initiation tab from the Restore Options dialog box, click **Schedule**, and enter your selections in the Schedule Details (Schedule Details) dialog box. Clicking **OK** from this dialog box saves your schedule. See Scheduled Data Recovery Operations for an overview of this feature.
   - If you want to run the job now, accept or click **Run Immediately** in the same tab and then click **OK**.

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

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

TIME RANGE OPTIONS FOR BASIC RESTORE

Before You Begin:

- Review the general and agent-specific restore requirements accessed from Restore Backup Data prior to performing any restore operation.

Required Capability: See Capabilities and Permitted Actions

To provide time range options for basic restore:

1. From the CommCell Browser, right-click the agent, instance or backup set that contains the data you want to restore, click **All Tasks** and then click the available **Restore** command (command names vary by agent).
2. In the **Restore** dialog box, type the starting path of the data you want to restore.
3. From the **Restore Options** and **Advanced Restore Options** dialog boxes, select the Time Range tab.
4. In the Advanced Restore Options (Time Range) dialog box, clear the **Show Deleted Items** option, if required.

   In the Specify Restore Time area select the **Time Zone** for the time range options. Select the **Exclude Data Before** option and select the start date and time from which you wish to restore data, and select the **Browse Data Before** option and select the end date and time until which you wish to restore data.
5. Click **OK** in the Advanced Restore Options dialog box.
6. Click **OK** in the Restore Options dialog box to execute the restore.

Back To Top

Browse and Restore

Topics | How To | Related Topics

OVERVIEW

There are two functions that help you retrieve backed up data from the backup media: browse and restore. In the CommCell Browser, the browse and variously-named restore commands appear, depending on agent, in the right-click menus at the agent, instance and/or backup set levels.

Browse operations allow you to view data that has been backed up by the agent on the client computer and select all or some of that data. Depending on the agent, there are several options available to customize your browse. See Browse Data for comprehensive information on Browse operations.

Restore operations allow you to retrieve data from backup media and restore it to the desired location. Restoring without browsing is most appropriate when you want to restore the latest backup job for an agent, instance or backup set and want to retain the current file structure. See Basic Restore for more information on restoring without using browse.

Browse and Restore

The Browse and Restore procedure is a sequential procedure that combines the two procedures. When you select a **Browse** command from the CommCell Browser, you can define and run one of many potential browse sequences. At the end of the browse, when you are looking at the resulting information presented in the Browse window, you can continue with a restore procedure simply by selecting data and clicking the **Recover All Data** button. As with the browse, depending on the agent, there are several options available to customize your restore.

Perform a **browse and restore** operation when you want to:

- restore from an earlier backup
- restore only select files/objects
- restore deleted files/objects
- when you don't want or don't need to retain the current file structure
- utilize browse options

Browse and Restore - How To

Topics | How To | Related Topics

BROWSE AND RESTORE

Before You Begin:

- Review the general and agent-specific restore requirements accessed from Restore Backup Data prior to performing any restore operation.

Required Capability: See Capabilities and Permitted Actions

To browse and restore data:

1. From the CommCell Browser, right-click the agent, instance, backup set, or Legal hold set (for Legal Hold data) that contains the data you want to restore, click **All Tasks** and then click the available **Browse** command (command names vary by agent).
2. Run a browse operation. See Browse Data for a list of customized browse operations and their step-by-step instructions. If you accept all defaults, you will be browsing the latest backups for the selected data.
3. From the Browse window, Select Objects From the Browse Window for Restore.
4. From the agent's **Restore Options** and **Advanced Restore Options** dialog boxes, select the restore options that you want to use. For agents with multiple tabs, do not click **OK** until you have used all of the desired tabs. When you accept all the default settings, you will be restoring the selected data to its original location. See Restore Backup Data for access to complete information on the agent-specific Restore Destination options and procedures available.
5. When restoring encrypted data, refer to Data Encryption for comprehensive feature information and procedures for using the Encryption tab of the Advanced Restore Options dialog box.
6. After completing your selections, you can either start an immediate restore or schedule the restore.
   - If you want to schedule the job, click the Job Initiation tab from the Restore Options dialog box, click **Schedule**, schedule the job, and then click **OK**.
   - If you want to run the job now, accept or click **Run Immediately** in the same tab and then click **OK**.

   While the job is running, you can right-click the job in the Job Controller and select **Detail** to view information on the job. After the data has been restored, you will see a job completion message in the Job Controller and Event Viewer.

Back To Top

Subclients - SAN *i*DataAgents

Topics | How To | Related Topics

OVERVIEW

The following table shows subclient creation and configuration details specific to SAN *i*DataAgents.

AGENT	Type of Data	Default Subclient created during install of the Agent	Supports Default Subclient	Supports User Defined Subclient	Contents of the default subclient when user-defined subclient is present	Other Types of subclients supported by the Agent	Notes
Image Level	volumes or mount points	No	No	Yes	N/A	CXBF (Unix), checksum (Unix)	None
Image Level ProxyHost	volumes or mount points	No	No	Yes	N/A	CXBF (Unix), checksum (Unix)	None
ProxyHost	volumes, folders, or mount points	No	No	Yes	N/A	None	None

For Image Level on Unix and Image Level ProxyHost on Unix, you can configure CXBF subclients or non-CXBF (checksum) subclients to enable snapshot and backup capabilities. For more information, see Snapshot Options.

CONFIGURABLE PROPERTIES

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

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

ACTIVITY CONTROL

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

CONTENT/DATABASES

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

DATA TRANSFER OPTIONS

Several configurable options to efficiently use available resources for transferring data secured by data protection operations are provided in the subclient. This includes the following:

- Enable or disable Data Compression either on the client or the MediaAgent.
- Configure the transfer of data in the network using the options for Network Bandwidth Throttling and Network Agents.

DATA ENCRYPTION

You can enable or disable the encryption of data for transmission over unsecure networks and for storage on media. For more information, see Data Encryption.

DATA PATHS

You can view the data paths associated with the primary storage policy copy of the selected storage policy or incremental storage policy. You can also modify the data paths for the subclient including their priority. For additional information, see Configuring Alternate Data Paths for Subclients.

DATA PROTECTION FILTERS

You can perform the following functions:

- Define data protection filters to exclude specified subclient data from being backed up or archived. For more information, see Filters.
- Use regular expressions (or wildcards) in subclient data protection exclusion filters. See Inclusions, Exclusions, and Exceptions to Exclusions for more information.
- Perform in-place editing of subclient data protection filter exclusions and exceptions. See Editing Filters for more information.

The Filters tab is only available for the ProxyHost *i*DataAgent.

PRE/POST PROCESSES

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

For Image Level or Image Level ProxyHost on Unix, if you configure a checksum subclient, you must provide the appropriate PreScan and PostBackup scripts to enable snap and mount capabilities. For more information on checksum subclients, see Snapshot Options. For more information on configuring the scripts, see Pre/Post Processes for Data Protection Operations: Image Level and Image Level ProxyHost *i*DataAgents.

SNAPSHOT OPTIONS

If installed on the client, QSnap can be enabled to back up locked files or to provide volume-level snapshot functionality and utilize the integrated block-filter driver. On Unix clients, for some supported agents, any volume that you add to a subclient is automatically configured as a CXBF Device, which is required to back up the volume. Other agents will require using Volume Explorer to create CXBF devices. Depending on the agent, and for specific scenarios, Volume Explorer can or should also be used to configure these devices. If QSnap is enabled, the CommCell Configuration Report displays a superscript Q in the subclient column. For step-by-step instructions for the supported file system *i*DataAgents, see Enable QSnap on a Subclient. Also, see Configure a CXBF Device in Volume Explorer as appropriate.

For Image Level and Image Level ProxyHost on Unix, you can configure subclients to enable a specified snapshot for backup. Specifically, you can configure a CXBF subclient to use QSnap, or you can configure a non-CXBF (checksum) subclient to use a supported snapshot that you desire. To this purpose, you can use the **Incremental Support Using** field in Subclient Properties (General). CXBF subclient configuration requires that QSnap be installed on the client; however, a QSnap install does not require you to configure the subclient as a CXBF subclient. If you use a snapshot other than QSnap (i.e., if you are configuring a checksum subclient), you must provide the appropriate PreScan and PostBackup scripts. For more information on configuring these scripts, see Pre/Post Processes for Data Protection Operations: Image Level and Image Level ProxyHost *i*DataAgents.

STORAGE POLICIES

You can associate the subclient to a storage policy. For more information, see Storage Policies.

SUBCLIENT NAME

You can rename a subclient. For step-by-step instructions, see Rename a Subclient.

USER ACCOUNTS

The following pertains to the ProxyHost, Image Level on Windows, Image Level ProxyHost *i*DataAgents:

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

See the section for your agent in User Accounts and Passwords for more information.

USER SECURITY

You can perform the following functions:

- Identify the user groups to which this CommCell object is associated.
- Associate this object with a user group.
- Disassociate this object from a user group.

For more information, see User Administration and Security.

VOLUME SHADOW SERVICE (VSS) BACKUPS

When using the Image Level *i*DataAgent in conjunction with a Windows *i*DataAgent that supports VSS, you can specify whether Volume Shadow Service (VSS) will be used to back up data for this subclient. If VSS is enabled, the CommCell Summary Report displays a superscript Q in the subclient column. For more

information, see VSS for the Image Level *i*DataAgent.

THINGS TO CONSIDER WHEN CREATING AND CONFIGURING SAN SUBCLIENTS

When creating and configuring subclients for SAN *i*DataAgents, keep in mind the following considerations:

IMAGE LEVEL AND IMAGE LEVEL PROXYHOST *I*DATAAGENTS

- Subclient content can only be either a volume or a mount point.
- The Windows 2000 and higher operating systems allow you to add new volumes to the existing file system name space without using new drive letters. For each volume that you add in this manner, Windows establishes a mount point, a pointer from the directory to the target data. Mount points are supported subclient content.
- For the Windows Image Level *i*DataAgent, you can add volumes to existing subclient content or remove them. However, whenever volumes are added or removed, the next backup job will be converted to a full backup.
- For Oracle BLI backups using the Image Level *i*DataAgent, each subclient is expected to quiesce one instance at a time. Do not configure an Image Level subclient's content with more than one Oracle instance; create separate subclients for each Oracle instance.
- For Solaris, if the slice 0 partition is empty and the disk is allocated space from slice 1 onwards, then slice 1 will be considered as slice 0 by the agent. So, from whatever slice partition you start to allocate space on the disk, that particular slice will be considered as slice 0.
- Do not specify the same volume as subclient content for both an Image Level and Image Level ProxyHost subclient.
- The size of a volume defined as subclient content will not be displayed until after the first backup and a refresh the CommCell Browser view.

PROXYHOST *I*DATAAGENT

- For the ProxyHost *i*DataAgent, subclients are used to back up different portions of the file system on a client computer. Each subclient is a designated subset of a production server to back up host data path mappings. You can create multiple subclients with each subclient containing a unique set of the production server data paths.
- There are certain situations in ProxyHost where running multiple simultaneous backups is not advisable. The snapshot utility can only be used by one subclient at a time, and any other subclient trying to access that utility while it is already in use will fail. Therefore, we advise that you stagger your multiple subclient backups so that the PreScan phase of each backup will not overlap.
- When you back up Exchange 2000 data with the ProxyHost *i*DataAgent, one storage group should not span multiple subclients.

Back to Top

Subclients - SAN *i*DataAgents - How To

Topics | How To | Related Topics

ADD/EDIT A DATA PROTECTION FILTER FOR A SUBCLIENT

Before You Begin

- Review Filters.
- Do not change the data protection or discovery filter of a subclient that has a data protection operation in progress.
- The system does not allow you to add entries that are not content of a particular subclient to that subclient's filter.
- For BlueArc and EMC Celerra (running at least DART OS 5.6.x) subclients, the filter string with or without wildcards must match the name or path of the file or directory being filtered.
- For NetApp subclients only name type filters (with wildcards) are supported:
  - You cannot enter paths as a filter (e.g. `/vol/vol0/data1`). Since NetApp does not support the use of paths in filters, if there are multiple files with the same name, even though they may be in different directories, all of them will be excluded from backups.
  - Name of the file or directory must exactly match the filter string.
  - You can specify a maximum of 32 strings in the exclude list.

Select the desired procedure:

- To add a data protection or discovery filter for a subclient
- To edit a data protection filter for a subclient

Required Capability: Capabilities and Permitted Actions

To add a data protection or discovery filter entry for a subclient:

1. From the CommCell Browser, right-click the subclient whose data protection or discovery filter you want to add, and then click **Properties** from the shortcut menu.
2. Click the Filters tab of the Subclient Properties dialog box.
3. For Exchange Mailbox, Exchange Mailbox/Public Folder Archiver Agents and Exchange Compliance Archiver, to specify a mailbox or folder that you want to exclude from data protection operations:
   - Click the upper **Add** button.
   - From the Browse window, expand the mailbox tree of the client computer.
   - Click the mailbox or folder that you want to exclude from the backup/archive operations on the selected subclient, and then click **Add**. (Repeat this step for each additional entry.)

- From the Browse window, click **OK**.

  The mailboxes or folders that you selected appear as entries in the upper pane. Repeat this step if you want to add more mailboxes and/or folders to the filter.

4. For NAS NDMP *i*DataAgents, click the **Add** button and, in the input window, type the name of the file, directory, (or path for BlueArc) that you want to exclude from the backups and click **OK**. The name displays as an entry in the **Exclude these files/folders/patterns** pane. Repeat this step if you want to add more files, directories, or paths to the filter.
5. For SharePoint Server *i*DataAgents, to specify a URL/file/folder/pattern that you want to exclude from data protection operations or to specify an exception filter for database backup sets, do one of the following:
   - Click the **Add** button next to **Exclude these files/folder/patterns**: input window and type the URLs of site collections that you want to exclude from the backups and click **OK**. The site collection displays as an entry in the **Exclude these files/folders/patterns** pane. Wildcards are supported. See Wildcards for more information. Repeat this step if you want to add more URLs, files, directories, or paths to the filter.
   - Click the **Add** button next to **Except for these files/folders**: input window and type the URLs/folders of site collections that you want to be exceptions for the exclusion filter and click **OK**. These exceptions will be included in the data protection operations. Wildcards are not supported for exception filters.
6. For Exchange Public Folder *i*DataAgents and SharePoint Server *i*DataAgent, to specify a workspace/folder that you want to exclude from the backups, click the upper **Browse** button and expand the *i*DataAgent of the client computer. Click the workspace/folder that you want to exclude from the backups and then click **Add**. Repeat this step for each additional entry.
7. For Windows/Unix/Macintosh File System *i*DataAgents, File Archiver for Windows/Unix Agents, and ProxyHost *i*DataAgent, to specify a file/folder/directory that you want to exclude from data protection operations, do one of the following:
   - Click the upper **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 the backups/archive operations. Repeat this step if you want to add more files/folders/directories to the filter.
   - Click the upper **Browse** button and expand the file system of the client computer. Click the file/folder/directory that you want to exclude from backups/archive operations and then click **Add**. Repeat this step for each additional entry.
8. For NetWare File System/NDS *i*DataAgents, to specify data that you want excluded from the backups, do one of the following:
   - To manually enter the path:
     - Click the **Add** button.
     - In the Input window, type the complete path (e.g., VOL1:\demo\disk2) of the data that you want to exclude.
     - Click **OK**. (The path that you typed displays as an entry in the upper pane.) Repeat this step if you want to exclude more data.
   - To browse and select a path:
     - Click the **Browse** button.
     - From the Backup Data window, expand the file system or NDS tree of the NetWare server.
     - Select the data that you want to exclude, and then click **Add**. Repeat this step for each additional entry.
     - From the Backup Data window, click **OK.** The data that you selected displays as entries in the upper pane.
9. For Lotus Notes Database and Lotus Notes Document *i*DataAgents, to specify a file or folder that you want to exclude from discovery, do one of the following:
   - Select the **Pattern/path to be excluded** field. Then type in a file or folder that you want to exclude. The format of the entry should start with a slash (\). The path entered is always relative to the data path of the partition. Then click **Add**. Repeat this step for each additional entry.
   - Click the upper **Browse** button and expand the file system of the client computer. Click the file or folder that you want to exclude from the backups and then click **Add**. Repeat this step for each additional entry.
10. For Exchange Mailbox, Exchange Mailbox/Public Folder Archiver Agents and Exchange Compliance Archiver, to specify a wildcard pattern of the folders that you want to exclude across all mailboxes within the subclient:
    - Click the lower **Add** button.
    - In the Input window, type the wildcard pattern of the folders that you want to exclude from backup/archive operations on the selected subclient.
    - Click **OK**. The path that you typed appears as an entry in the lower pane. Repeat this step if you want to add more entries to the filter.
11. For Exchange Public Folder *i*DataAgents and SharePoint Server *i*DataAgent, to specify an exception to an excluded workspace/folder (i.e., a folder/document that you want included in the backups, but whose parent directory has been excluded), click the lower **Browse** button and expand the *i*DataAgent of the client computer. Click the workspace/folder that you want to include in the backups and then click **Add**. Repeat this step for each additional entry.
12. For Windows/Unix/Macintosh File System *i*DataAgents, File Archiver for Windows/Unix Agents, and ProxyHost *i*DataAgent, to specify an exception to an excluded folder/directory (i.e., a file or folder/directory that you want included in the data protection operations, but whose parent folder/directory has been excluded), do one of the following:
    - Click the lower **Add** button and, in the **Enter Path** window, type the complete path (including drive letter) of the file/folder/directory that you want to include in the backups/archive operations. Repeat this step if you want to add more exceptions to the filter.
    - Click the lower **Browse** button and expand the file system of the client computer. Click the file/folder/directory that you want to include in the

backups/archive operations and then click **Add**. Repeat this step for each additional entry.

13. For NetWare File System/NDS *i*DataAgents, to specify an exception to an excluded directory or NDS container (i.e., data that you want included in the backups, but whose parent directory or NDS container has been excluded), do one of the following:
    - To manually enter the path:
      - Click the **Add** button.
      - In the Input window, type the complete path (e.g., VOL1:\demo\disk2\readme) of the data that you want to include.
      - Click **OK**.
    - To browse and select a path:
      - Click the **Browse** button.
      - From the Backup Data window, expand the file system or NDS tree of the NetWare server.
      - Click the file or folder that you want to include in the backups and then click **Add**. Repeat this step for each additional entry.
      - Click **OK.** The selected data displays as entries in the lower pane.
14. Click **OK** to save your changes.

Required Capability: Capabilities and Permitted Actions

To edit a data protection filter entry for a subclient:

1. From the CommCell Browser, right-click the subclient whose data protection filter you want to edit, and then click **Properties** from the shortcut menu.
2. Click the Filters tab of the Subclient Properties dialog box.
3. Click the filter entry that you want to edit, and then click the **Edit** button associated with that pane.
4. Type the changes into the **Enter Path** dialog box, then click **OK**.
5. Click **OK** to save your changes.

NOTES

- When you change a data protection or discovery filter, the change is effective the next time a data protection operation is run on the applicable subclient.
- Performing a full backup after changing filters or exceptions is recommended.

ASSOCIATE A SUBCLIENT TO A STORAGE POLICY

Required Capability: See Capabilities and Permitted Actions

To associate a subclient to a storage policy:

1. From the CommCell Browser, right-click the subclient whose associated storage policy you want to change, then click **Properties** from the shortcut menu.
2. Click the Storage Device tab of the Subclient Properties dialog box.
3. From the **Storage Policy** list of the **Data Storage Policy** tab, select a data storage policy to associate with this subclient. If necessary, click the **Create Storage Policy** button to create a new storage policy to which the subclient can then be associated.
4. From the Changing a Storage Policy window select the next type of backup operation. Click **OK**.
5. If applicable for your agent, you can change the number of data streams from the **Number of Data/Database Backup Streams** field.
6. If applicable for your agent, click the **Log Storage Policy** tab and select a storage policy to associate with this transaction log subclient from the **Transaction Log Storage Policy** list. Also, you can set the **Number of Transaction Log Backup Streams** from this tab.
7. Click **OK** to save your changes and close the Subclient Properties Storage Device tab.

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

Required Capability: See Capabilities and Permitted Actions

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

1. From the CommCell Browser, expand the tree to view the appropriate level icon for the affected agent.
   - From the agent, instance/partition, or backup set/archive set level, right-click the appropriate icon, click **All Tasks**, and click **New Subclient** from the short-cut menu.
   - From the subclient level, right-click the subclient icon and click **Properties** from the short-cut menu.
2. From the Subclient Properties dialog box, create and/or configure the subclient as appropriate. Then click the **Pre/Post Process** tab.

3. From the **Pre/Post Process** tab, click **Change**.
4. From the User Account dialog box, select one of the account options. If you select **Impersonate User**, type the appropriate user name and password.
5. Click **OK** to save the settings.

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

Before You Begin

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

Required Capability: Capabilities and Permitted Actions

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

1. From the CommCell Browser, right-click the subclient for which you want to configure a pre/post process, and then click **Properties** from the shortcut menu.
2. Click the Pre/Post Process tab of the Properties dialog box.
3. For an agent other than the Oracle RAC *i*DataAgent, click inside the space that corresponds to one of the following phases and type the full path of the process that you want executed during that phase. Alternatively, click **Browse** to locate the process (applicable only for paths that do not contain any spaces). For the Oracle RAC *i*DataAgent, click **Browse** for the corresponding process, click the name of the control node client in the Select Client for Browse dialog box, and click **OK**. Then browse for and click the process.
   - PreBackup
   - PreScan
   - PreArchive
   - PreCopy
   - PreSnap
   - PostBackup
   - PostScan
   - PostArchive
   - PostCopy
   - PostSnap

   Click **OK**.
4. If you want to run a Post Process for all attempts to run that job phase, then select the corresponding checkbox.
5. For subclients on Windows platforms, if **Run As** displays **Not Selected**, or if you want to change the account that has permission to run these commands, click **Change**.
   a. In the User Account dialog box, select **Use Local System Account,** or select **Impersonate User** and enter a user name and password. Click **OK**.
   b. If you selected Local System Account, click **OK** to the message advising you that commands using this account have rights to access all data on the client computer.
6. Click **OK** to save your changes and close the Pre/Post Process tab of the Properties dialog box.

CONFIGURE SUBCLIENT CONTENT

Before You Begin

- Review Subclients.
- Do not configure the content of a subclient while the parent node or any sibling subclient has a data protection or archive operation currently running on it.
- Exchange Mailbox *i*DataAgents and Exchange Mailbox/Public Folder Archiver Agents: If you change the contents of the default backup set or archive set then the auto-discover feature will be disabled. If you disable the auto-discovery feature, newly created mailboxes will not be backed up/archived unless they are manually discovered and assigned to a subclient.
- NAS NDMP *i*DataAgents: You must ensure there is no overlap in content between all subclients. Overlap in subclient content will result in loss of data. An existing subclient's contents are not automatically changed when another subclient is added with overlapping contents.
- SharePoint Server *i*DataAgent: The Site Content Database, the Site Collection Database, the Site Database, and the Site Index for the virtual server must all

be assigned to the same subclient.

- Lotus Notes Document *i*DataAgent: Review Assigning Restore View Names to Newly-discovered Databases
- QR Agent: Follow these guidelines when adding a volume to a QR Agent subclient:
  - The volume must correspond to a physical disk or RAID array.
  - A volume created by volume management software other than VxVM is not valid subclient content.
  - Subclients may have overlapping content; however, if two or more subclients overlap, they all must use the same snapshot engine. If the QR policies associated with the subclients are configured to use different snap engines, they must be reconfigured to use the same snap engine in this scenario.
- **Caution Against Re-configuring Default Subclient Content**

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

Required Capability: See Capabilities and Permitted Actions

To configure subclient content:

1. From the CommCell Browser, right-click the subclient for which you want to configure content, click **All Tasks** (if applicable) and then click **Properties**.
2. Follow the procedure below that is applicable for your agent:
   - For File System, Active Directory, File Archiver, Exchange Public Folder *i*DataAgents, NDS, and SharePoint Server *i*DataAgents click the Subclient Properties (Content) tab and configure content for the subclient as described below for your agent:
     - For File System, Active Directory, File Archiver, NDS, and SharePoint Server *i*DataAgents: Type the full path of the data that you want to include as subclient content in the **Enter New Content** field, then click **Add**. Optionally, click **Browse** to enter the content. When browsing content while configuring SharePoint subclients, you can add content via multiple selections with the CTRL or SHIFT keys. For Windows, when specifying a UNC Path, click **As User**, and enter the user account information for the domain user with permissions for that path. For NetWare/DNS, see the Notes section below for content path examples. For Unix File Systems, you can enter the mount point of an NFS-mounted file system, see the Notes section below for examples.
     - For Exchange Public Folder *i*DataAgents: Click **Browse**, select folders to include as content, then click **Add**.
     - For the Unix File System *i*DataAgents, to facilitate the management of resource fork data in Apple double-encoded Macintosh files, click **Enable Apple Double Support**.
     - For the Unix File System *i*DataAgents, to view the actual data path for any symbolic link in the subclient content, click **Expand symbolic links of subclient content** and then click **Discover**.
   - For NAS NDMP *i*DataAgents, configure the **Backup Content Path** field(s) as described below, then click **Add**:
     - Click the drop-down list arrow to display the root volumes on the file server. To change the root volume, click one in the list. If you want to refine the content path further, use the space to the right of (or below) the root volume list to enter additional path information. Note the following:
       - For NetApp, the root volume is the mount path of each volume.
         Example: for volume `FS1` the root volume will be `/vol/FS1`.
       - For EMC Celerra, the root volume is the mount point created for a volume.
         Example: for volume `FS1` with mount point `/FS1` the root volume will be `/FS1`.
       - For Hitachi, no root volumes are shown in the drop down list. Type the full path of the root volume.
         Example: for volume `FS1` with mount point `/mnt/FS1` the root volume will be `/mnt/FS1`.
       - For BlueArc, the root volume is a combination of a descriptor of the path and the volume name.
         Example: for volume `FS1` with a mount point of `/` the root volume will be `/__VOLUME__/FS1`.
     - Optionally, for NetApp NAS NDMP, click **Browse** to enter the content.
   - For Exchange Mailbox and Exchange Mailbox/Public Folder Archiver Agents follow the procedure to Discover and Assign New Mailboxes or Assign Mailboxes to Another Subclient.
   - For Lotus Notes Database and Document *i*DataAgents follow the procedure to Discover and Assign New Databases or Assign Databases to a Subclient.
   - For DB2, DB2 DPF, Exchange Database, Novell GroupWise, SharePoint Server, SQL Server Database, Sybase, and MySQL *i*DataAgents, click the Subclient Properties (Content) tab and configure content for the subclient as described below for your agent:
     - For the DB2 *i*DataAgent, specify whether you want to include the entire database or a subset of this data as content for the subclient. For the DB2 DPF *i*DataAgent, specify whether you want to include all the affected database partitions or a subset of this data as content for the subclient.
     - For Exchange and GroupWise *i*DataAgents: Click **Configure**. From the Add/Modify Subclients dialog box click the subclient entry for the database element/Storage Group that you want to add to the new subclient and select the name of the destination subclient from the list that appears. Alternatively, you can select and assign a range of databases/storage groups using the **Change all selected databases/storage groups to** list. Note that you must have at least one database element/Storage Group assigned to this subclient in order to save the configuration.

       A database/Storage Group that is not configured for a subclient does not appear in the list. This can be the case if the subclient containing the database/Storage Group was deleted. If this happens, click **Discover** to display all databases/Storage Groups.
     - For the SharePoint Server *i*DataAgent, follow the procedure to Discover and Assign New Data Types.
     - For the Sybase *i*DataAgent, follow the procedure to Manually Discover Databases.

- For the MySQL *i*DataAgent, follow the procedure to Configure MySQL Databases.

- For the Informix *i*DataAgent, click the Subclient Properties (Content) tab and define the contents of the subclient. Specifically, establish the backup mode for the data to be backed up, set the backup level, and decide whether to back up the emergency boot file and/or the ONCONFIG file.
- For the Oracle, SAP for Oracle, or Oracle RAC *i*DataAgent, click the Subclient Properties (Content) tab and define the contents of the subclient. To configure this subclient for specialized types of backups, follow the appropriate procedure below:
  - Create Subclient for Backing Up Archived Redo Log Files
  - Create Subclient for Backing Up Offline Databases
  - Create Subclient for Backing Up Online Databases
  - Create Subclient for Performing Selective Online Full Backups
- For SAN *i*DataAgents, click the Subclient Properties (Content) tab and configure content for the subclient as described below for your agent:
  - Image Level on Unix *i*DataAgent: Click **Add**. From the Add Content Path dialog box, select the volume(s) that you want to back up (use CTRL + click to select multiple volumes). Click **OK**. The selected volumes are added to the **Contents of subclient** list. These volumes are automatically configured to be CXBF devices. Alternatively, use Volume Explorer per specific scenarios to configure CXBF devices.

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

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

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

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

    See also Configure Subclients for Overlapping Content.

3. Click **OK** to save your content configuration.

NOTES

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

  <primary_host_path> --> <backup_host_path>

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

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

CONFIGURE THE SUBCLIENT FOR DATA ENCRYPTION

Encryption settings made at the subclient level are for data protection and recovery operations run from the CommCell Console and are not related in any way to settings made at the instance level which is for third-party Command Line operations only.

See Data Encryption - Support for a list of supported products.

Before You Begin

- Encryption must be enabled at the client level prior to configuring any subclients residing on that client. See Configure the Client for Data Encryption.
- If you are attempting to configure for third-party Command Line operations, do not use this procedure. See Configure Third-party Command Line Operations for Encryption.

Required Capability: Capabilities and Permitted Actions

To configure the subclient for data encryption:

1. From the CommCell Console, right-click the subclient and click **Properties**.
2. From the Subclient Properties (Encryption) tab, select an option based on the criteria described in the Encryption tab help.
3. Click **OK** to save your settings and close subclient properties.

CREATE A NEW SUBCLIENT

Before You Begin

- Review Subclients.
- Do not create a subclient while the parent node or any sibling subclient has a data protection or archive operation currently running on it.
- In cases where a new subclient is created with the same name as a deleted subclient, the system will append a Unix time stamp to the deleted subclient's name in data protection job history reports and views to distinguish the two subclients. For example, *subclientname*_1104257351.
- Informix *i*DataAgents: If you will be using the Informix ONBAR utility to create backup and restore scripts, you need not create subclients. Otherwise, if you will be using the CommCell Console to back up and restore Informix database objects (subsets/dbspaces), then you will need to create a subclient.
- ProxyHost *i*DataAgents: If you are using a BCV, you must prepare a batch file or a shell script file on the backup host containing commands to synchronize and split the BCV. The Resource Pack includes information on configurations for these batch files or shell scripts, as well as examples that apply to specific applications and hardware (e.g., Exchange databases in an EMC Symmetrix environment). See Resource Pack for more information on the Resource Pack.

  The ProxyHost *i*DataAgent also requires that you set permissions for the batch/shell script file on the backup host.
- SQL Server Database *i*DataAgents: When running on Windows Server 2003 and VSS is enabled, the **New Subclient** command is not available.
- PostgreSQL *i*DataAgents: Once you configure the PostgreSQL instance, the system automatically generates the default backup sets and default subclients. However, you can use the CommCell Console to create user-defined subclients for dump backup sets to distribute some of the database content. You cannot create user-defined subclients for FS backup sets.

Required Capability: See Capabilities and Permitted Actions

To create a new subclient:

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

- For the Windows *i*DataAgents that support VSS, you can optionally Enable VSS on a Subclient.

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

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

This task is now complete.

DELETE A DATA PROTECTION FILTER FROM A SUBCLIENT

Before You Begin

- Review Filters.
- Do not change the data protection or discovery filter of a subclient that has a data protection operation in progress.
- File Archiver for Windows/Unix Agents: We recommend that you don't delete the following entries from the exclusion filter, as it could cause your file system to be inoperable. For Windows, these include: *.dll, *.bat, *.exe, *.cur, *.ico, *.lnk. For Unix, these include *.a, *.ksh, *.csh, *.sh, *.lib, *.so.

Required Capability: Capabilities and Permitted Actions

To delete a data protection or discovery filter entry from a subclient:

1. From the CommCell Browser, right-click the subclient whose data protection or discovery filter you want to delete, and then click **Properties** from the shortcut menu.
2. Click the Filters tab of the Subclient Properties dialog box.
3. To delete an entry from the Exclusions list, click the entry in the upper pane then click the upper **Delete** button. (Repeat this step for each entry that you want to delete.)
4. To delete an entry from the Exceptions list (if applicable for your agent), click the entry in the lower pane then click the lower **Delete** button. (Repeat this step for each entry that you want to delete.)
5. Click **OK** to save your changes.

NOTES

- Whenever you delete an entry from the exclusion filter, check if the exceptions list (i.e., lower pane) contains any entries that are children of the deleted data (if applicable for your agent). If so, you should delete them as described in Step 4 since they no longer need to be listed. The system automatically deletes any exceptions that are children of a deleted exclusion unless you used wildcard expressions in the exclusion path.
- When you change a data protection or discovery filter, the change is effective the next time the subclient is backed up/archived.
- Data will not be backed up in a differential backup for a subclient after a filter was removed.
- Since Incremental backups only back up data that has been modified since the last backup, previously filtered files whose filters are now removed, will not be backed up unless they have been modified since that last backup. To back up previously filtered files that have not been modified but whose filters have been removed since the last backup, you need to run a Full backup.
- Performing a full backup after changing filters or exceptions is recommended.

DELETE A USER-DEFINED SUBCLIENT

Related Topics:

- Command Line Interface - qdelete subclient
- Subclients

Required Capability: See Capabilities and Permitted Actions

To delete a user-defined subclient:

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

ENABLE OR DISABLE OPERATIONS

Required Capability: See Capabilities and Permitted Actions

Level	Capability

CommCell	Administrative Management with CommCell level association
Client Computer Group	Administrative Management with Client Computer Group level association
Client	Agent Management with Client level association
Agent	Agent Management with Agent level association
Subclient	Agent Management with Subclient level association

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

1. From the CommCell Browser, right-click the CommServe, client computer group, client computer, agent, or subclient, and then click **Properties** from the short-cut menu.
2. From the Activity Control tab of the associated Properties dialog box, select or clear option(s), as desired.
3. Click **OK**.

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

Enable or Disable Software Compression for a Subclient

Before you Begin

- Do not enable/disable software compression for a subclient that is being backed up/archived.

Required Capability: Capabilities and Permitted Actions

To enable software compression for a subclient:

1. From the CommCell Browser, right-click the subclient for which you wish to enable software compression and then click **Properties**.
2. Click the **Storage Device** tab and from the Data Storage Policy tab, select the storage policy from the **Storage Policy** list.

   If applicable for the selected agent, click the Log Storage Policy tab and select a storage policy from the **Transaction Log Storage Policy** list.
3. Then click the Storage Device (Data Transfer Option) tab and choose the appropriate compression option for this subclient.
4. Click **OK** to save your changes.

This task is now complete.

Remove a Process from Pre/Post Processing of Data Protection/Archive Operations

Before You Begin

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

Required Capability: Capabilities and Permitted Actions

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

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

RENAME A SUBCLIENT

Before You Begin

- You can rename a subclient at any time. However, we recommend that you don't rename a subclient while a data protection or archive operation is running on that subclient.
- In cases where a subclient is renamed using the same name as a deleted subclient, the system will append a Unix time stamp to the deleted subclient's name in data protection job history reports and views to distinguish the two subclients. For example, *subclientname*_1104257351.

Required Capability: See Capabilities and Permitted Actions

To rename a subclient:

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

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

SET THE NETWORK BANDWIDTH AND NETWORK AGENTS FOR A DATA PROTECTION OPERATION

Before you Begin

- Do not modify the network bandwidth and network agents for a subclient or instance that is being backed up.

Required Capability: Capabilities and Permitted Actions

To Set the Network Bandwidth and Network Agents for a Data Protection Operation:

1. From the CommCell Browser, right-click a subclient and then click **Properties.**

   For the DB2, DB2 DPF, Informix, Oracle, Oracle RAC, SAP, or Sybase *i*DataAgent, right-click an instance and then click **Properties.**
2. Click the **Storage Device** Data Transfer Option tab.

   For the QR Agent:

   - To control network bandwidth settings, use the Throttle Network Bandwidth section in the General tab of the Subclient Properties dialog box.
   - To control the number of network agents, you must create a `nQRNetworkAgents` registry key.
3. Enter a number of **Network Agents** that must be used to perform data protection operations on the subclient/instance.
4. Click the **Throttle Network Bandwidth (MB/HR)** option and then enter the throughput as needed. Note that throttling is done on a per Network Agent basis.
5. Click **OK** to save the changes.

This task is now complete.

VIEW DATA PATHS ASSOCIATED WITH A SUBCLIENT

Required Capability: See Capabilities and Permitted Actions

To view data paths:

1. From the CommCell Browser, right-click the subclient whose data paths you want to view, then click **Properties** from the shortcut menu.
2. Click the Storage Device tab of the Subclient Properties dialog box.
3. From the **Data [or Logs] Storage Policy** tab, click **Show Data Paths** to view the data paths used by the subclient to access the storage media for data protection operations. Click **Close** to exit the Data Paths dialog box.
4. Click **OK** to exit the Subclient Properties Storage Device tab.

VIEW SUBCLIENT CONTENT

Required Capability: See Capabilities and Permitted Actions

To view content of a subclient:

1. From the CommCell Browser, right-click the subclient whose content you want to view, then click **Properties**.

2. From the Subclient Properties dialog box, click the **Content** tab (or **Databases** tab for Lotus Notes) to view the contents of the subclient.
3. Click **OK** to close the dialog box.

Back To Top

Overview - The QSnap® Service

Topics | How To | Related Topics

INTRODUCTION

The QSnap® Service is a software-based snapshot implementation that integrates with other Agents, providing all of the components necessary for basic snapshot functionality without requiring specialized hardware. The QSnap product is an installable, licensed software module. In this release, the QSnap service will still be automatically installed with CDR but will not be used as the default snap engine. QSnap® Services will be used as the source snapshot on Windows 2000 and Windows XP. It will not be used for creating Recovery Points, even with the aforementioned operations systems.

On Unix platforms, the driver name is Unix QSnap. You will often see "CXBF" in references to the Unix QSnap driver. For example, a CXBF device is a volume or partition that is monitored by the CXBF block-filter driver.

The QSnap service provides the following functionality:

- **Quick Recovery and Data Replication Agents**
  QSnap functionality is automatically installed with the Quick Recovery and ContinuousDataReplicator agents. In addition to providing volume-level snapshot functionality, the QSnap driver has an integrated block-filter driver that monitors changes to a volume over time. For any volume-level agent, the bitmaps help ensure that the next data protection operation can be an incremental backup instead of a full backup.
- **Unix File System *i*DataAgent**
  QSnap services can be used with the Unix File System *i*DataAgent. When QSnap is installed and enabled for use with this Agent, the snapshot functionality of QSnap is used to snap and then back up locked and/or open files.
- **Windows File System *i*DataAgent**
  QSnap services can be used with the Windows File System *i*DataAgent. When QSnap is installed and enabled for use with this Agent, the snapshot functionality of QSnap is used to snap and then back up locked and/or open files.
- **Backup & Recovery Agents**
  QSnap services is automatically installed with the following Backup & Recovery *i*DataAgents: Image Level, and Image Level ProxyHost.

Back to Top

CONFIGURATION

For all Agents, the QSnap service must be installed prior to configuring the software. If your Agent supports integration with the QSnap service, an install procedure that includes the steps for installing QSnap software has been provided. See Installation for installation procedures.

QSnap configuration differs depending upon the Agent with which it is working; select your Agent from the following list for specific configuration information and procedures:

- QSnap for ContinuousDataReplicator
- QSnap for the Image Level *i*DataAgent
- QSnap for the Image Level ProxyHost *i*DataAgent
- QSnap for the Quick Recovery Agent (includes Recovery Director)

- QSnap for the Unix File System *i*DataAgent
- QSnap for the Windows File System *i*DataAgent

Back to Top

QSNAP® COPY-ON-WRITE CACHE

The QSnap creates a copy-on-write (COW) cache, to which it copies blocks that are being overwritten on the source volume during a snapshot creation. This preserves the original data, ensuring an accurate snapshot for data protection and recovery operations.

The COW cache contains only copies of blocks that have been overwritten. Any new data that is written to free space on a source volume is not cached. It is important to ensure there is enough space for the cached blocks.

By default on Windows platforms, the COW cache minimum size is 50MB, and the system expands the cache file as needed **up to 90% of the total capacity of the volume**. A snapshot will be terminated if it causes the cache file to go beyond this size.

You can override the default cache sizes by entering megabyte values for minimum cache and maximum cache in the General tab of the *i*DataAgent's Properties window. (See Change the COW Cache Size.)

But 90% of a volume's total capacity for maximum is a hard upper limit. Although the system will accept a megabyte value that represents more than this limit (e.g., 95 MB on a 100 MB volume), the snapshot will still terminate when the cache reaches 90% of total volume capacity (in the example: 90 MB).

Meanwhile, if you want to reset the minimum cache, note that 25 MB is the lower limit. If you try to set the minimum below this limit, the software will consider the limit to be 25 MB. And, when resetting the maximum and minimum values, be sure not to set the minimum value higher than the maximum value or limit.

Another factor is whether the data is written to free space on a volume. For example, if 550MB of new data is written to free space on a source volume, no blocks are overwritten and therefore no blocks are cached, and the default COW cache settings do not require adjustment. The snapshot driver does not need to cache data written to free space in order to create a snapshot. However, if 550MB of data is written over existing data and the 550MB represents more than 90% of total capacity, then the default maximum cache size will be exceeded and the snapshot will terminate.

Therefore, when determining the maximum size of the copy-on-write cache, consider the following, in order of importance:

- Whether the I/O will write to free space, or change existing blocks
- Amount of I/O
- Size of the source volume
- The number of files on the source volume

Be sure to set the maximum cache high enough (not to exceed 90% of the volume's total capacity) to accommodate these factors.

> The COW cache must be located on a volume with a supported file system.

WINDOWS COW CACHE

By default, a copy-on-write (COW) cache is created on each source volume. Whether you are using NTFS or FAT/FAT32 for the source volume, any NTFS volume is supported for the COW cache location, with the following caveats:

- The cache partition cannot be on a network share, a mapped drive, or a mount point.
- Not all Agents support all source volume types and file systems. See the Supported Data Types in the Product Overview for your Agent.

CHANGING THE CACHE LOCATION ON WINDOWS

On the Windows platform, the COW cache location can be changed through the Client level properties. Note that if you specify an alternate COW cache location, all of the COW caches for all of the volumes on the client will use that COW cache location. Be sure there is enough free space to account for the caches.

If you are using the QSnap services with the Windows File System iDataAgent in a **cluster environment**, and you want to change the COW cache location:

- On a physical machine, do not set the alternate cache location to a virtual machine resource. Use only a non-shared drive on the physical machine.
- On a virtual machine, do not set the alternate cache location to a physical machine resource. Use only a shared virtual drive.

See Change the COW Cache Location for step-by-step instructions.

CACHE SIZE ON WINDOWS

The COW cache size is adjustable through the Agent level properties. You can set the Minimum and Maximum cache sizes for the volumes associated with the Agent. By default, the minimum size is 50MB and the maximum size is 90% of the total capacity of the volume.

See Change the COW Cache Size for step-by-step instructions.

CHANGING THE WRITE-INACTIVITY-PERIOD

The QSnap enabler requires a short period of no disk activity to create the snapshot. You may require a longer write inactivity period due to slow disk

performance. See Change the Write Inactivity Period (WIP) for step-by-step instructions.

CACHE LOCATION FOR QUICK RECOVERY AGENT SUBCLIENT USING RECOVERY POINTS

If you are creating a subclient that will use a QR™ Policy with QSnap as the snapshot enabler and Recovery Points enabled, you must specify a cache volume (unless a cache volume had been specified in Client Properties/Advanced). The cache volume specified during Quick Recovery® Agent subclient creation will be used as the cache volume for all snapshots (source and destination) and for all Agents for that particular machine.

CONSIDERATIONS - QSNAP COW CACHE FOR WINDOWS

Consider the following about the COW cache size, location, configuration, and use:

- **Snapshot Cache Cannot be used for Copyback** Do **not** select the snapshot cache volume as the copyback destination even though it will show up as a possible copyback destination if it is larger than the source volume. During copyback operations, the destination volume needs to be unmounted, but a snapshot cache volume must remain mounted during the entire copyback operation. Selecting the cache volume as the copyback destination will cause the job to fail.
- **Cache Considerations for ContinuousDataReplicator**

  CDR creates non-persistent snapshots during the Baselining phase for Data Replication, on Windows 2000 and Windows XP computers. The following information should be considered in the context of disk space requirements for your particular environment and configuration.

  - **Source Computer:**
    - The default cache location for the snapshot created during the SmartSync Scan phase is on the source drive specified for each Replication Pair.
    - If you specify an alternate COW cache location in the Client Properties of a ***source*** computer, that location will be used whenever a snapshot is created during the SmartSync Scan phase for all Replication Pairs in all Replication Sets on that client computer.
    - The default minimum cache size is 50MB. This can be increased using the `nIncrementalCache` registry key, and is recommended for faster computers if sufficient space is available.

  To configure an alternate cache location, see Changing the Cache Location on Windows above.

- **Cache Volume becomes full**

  When a cache volume becomes full (reaches its maximum allowed size) any current snapshot is terminated and all snapshots for that source volume are deleted. Even though the snapshots have been deleted, the Recovery Point entries remain in the system database, and thus appear during a Browse operation; however, if you attempt to mount, unmount, copy back, or delete those Recovery Points, you will receive an error message. At this point, you can check the Windows Event Viewer for an event that indicates all snapshots were deleted for a particular volume, indicating this is what happened; if so, you should delete those Recovery Points in the Browse window.

  When using VSS or QSnap enabler with either the Quick Recovery Agent or ContinuousDataReplicator, it is recommended that you set the `Disk Space Low` alert to notify you when disk space is getting low; this can help you diagnose lack of cache space as a problem (or possibly avoid it). When configuring this Alert, include the <DISK SPACE INFO> Alert Token in the email notification, so that meaningful data is included in the emailed alert message. Also, consider configuring the nMACHINE_MAINT_INTERVAL_MINUTES registry key if the Space Check should be performed more often than the default of 24 hours. For more information, refer to Space Check for the Quick Recovery and ContinuousDataReplicator Agents.

- **Using RAW Volumes**

  To use a RAW volume as the QSnap COW Cache volume, the value of the registry key QSCacheVolume must be created. This value designates which RAW volume will be used for the cache. Furthermore, before using a RAW volume, the cache volume type must be set with the QSCacheVolumeType key. If cache values are manually set in the registry, they will take precedence over the values set in the CommCell Console.

UNIX COW CACHE

For Unix, the mount point for the copy-on-write (COW) cache is specified during installation. For instructions on mounting a partition to the COW cache mount point, see Mount the COW Cache Partition. It is recommended that you use the local file system for COW cache.

The COW cache must be located on a volume with a supported file system.

The cache partition cannot be on a network share.

After installation, you can assign any volume/partition to this mount point so long as it is one of the supported configurations below:

Source Volume	Supported COW Cache Location	Notes
JFS2 - AIX	• Any volume other than the source volume	• Volumes managed by VxVM can be used for the COW cache partition; however, VxFS volumes cannot be used. • Only one COW cache partition is allowed per client. It is recommended that this partition be a dedicated cache partition (not used for other purposes). • For the QR Agent on Solaris, the COW cache partition must be on a CXBF device. • On UNIX, the source volume is not supported as the COW cache partition. You must specify a
EXT2, EXT3, EXT4 - Linux	• Any volume other than the source volume	
XFS - Linux	• Any volume other than the source volume	
Reiserfs - Linux	• Any volume other than the source volume	
UFS - Solaris	• Any volume other than the source volume, with a UFS file system	
VxFS - Solaris	• Any volume other than the source volume, with a UFS file system	

		volume other than the source.

New data does not use a Unix COW cache on Solaris volumes or on Linux ext2/ext3 volumes; only *changed* data on these volumes use the cache. However, both *new* data and *changed* data use the Unix COW cache on AIX volumes or Linux volumes using a file system other than ext2/ext3.

Not all agents support all source volume types and file systems. See Supported Data Types in the Product Overview for your agent.

CALCULATE COW CACHE SIZE ON UNIX

The maximum size of the COW cache equals the size of the volume mounted to the COW cache mount point (which was specified during the software installation). If the COW cache size exceeds the size of this partition, then a larger partition must be mounted to this mount point.

The size of the cache partition should be determined as follows, based on the assumption that 25% of a volume is being modified between the times of snap-on to snap-off on each volume:

If m is the number of volumes for which the cache is being created with size s, and n is the number of snapshots that are expected to be copied, then the cache size should be:

(n x s)/4

For example, if there are 20 Volumes (m) of 10GB size (s) and we expect to run snapshots on 2 volumes (n) at a time then the cache partition size should be at least:

(2 x 10)/4= 5GB

COW CACHE FOR LINUX VOLUME GREATER THEN 3TB

Reformat the file system containing the CXBF cache to accept sparse files 3TB or bigger (XFS, EXT4, Reiserfs) in size.

Set the value of dSnapChunkSize registry key to 1048576.

COW CACHE FOR RED HAT LINUX

The Red Hat Linux file system chooses a cache block size according to disk size. So if you are using a small disk for cache, the file system will use a small block size such as one kilobyte. A small block size and a source disk with many files can cause the CXBF driver scan to fail. For this situation, it is recommended that you manually specify a block size of 4k (4096bytes) with the following command: `mke2fs -b 4096`.

CALCULATE BITMAP SIZE ON UNIX

For the cache size, we use one bit per extent to mark if it is modified or not. The default extent size is 2MB.

547 GB = 280,064 extents

280,064 extents needs 35,008 bytes

So, for 547GB, we need only 35 KB of space in /etc/galaxy directory.

The cache size can be as large as the volume. A 547 GB file system completely modified from scratch will occupy 547GB of space. Also, please note that this is a sparse file. The size shown by ls -l is not an indication of the space consumed. "du -sk" on the file is the correct measure on the actual blocks used.

Back to Top

THE BLOCK-FILTER DRIVER AND BITMAPS

Through the use of a block-filter driver, the QSnap software creates bitmaps to track the block-level changes to a volume over time. These bitmaps are stored in memory and later written to disk. The bitmaps help ensure that the next data protection operation can be an incremental backup instead of a full backup. Unlike the COW cache, which is created and used only when creating snapshots, bitmaps are always maintained for devices monitored by the block-filter driver.

BLOCK-FILTER ACTIVATION

When the QSnap software is installed, its block-filter is **activated** for all volumes by default.

MANUALLY ACTIVATING/DEACTIVATING THE BLOCK-FILTER

The block-filter may be active on volumes you do not want it to be active on — as may be the case after upgrading QSnap. You can manually activate/deactivate the QSnap block-filter driver on specific volumes, using the following:

- Qsnp2Config - See Enable/Disable QSnap Block-filter Driver on Specific Volumes for step-by-step instructions.
- QST2 - See QST2 Tool for step-by-step instructions.
- BFActivateAll registry key - An alternative method to enable or disable the block filter of all volumes on the machine.

BITMAPS AND CHANGED BLOCKS

After the first full backup, you can update the backup incrementally so that only the changed blocks on the source volume are backed up or copied. In order to

keep track of the changes, the QSnap program creates a bitmap that records the changed blocks for each source volume.

If the system cannot verify the integrity of a bitmap, it will force the next backup to be a full backup, which may be undesirable in the following situations:

- moving very large amounts of data
- moving data across a slow LAN or WAN
- (for the Quick Recovery® Agent) recreating off-site QR™ recovery volumes

---

PERSISTENCE

To avoid unwanted full backups or QR™ Volume Creation jobs, you can configure QSnap® bitmap persistence on a volume, to capture a bitmap of the blocks that have not been read and copied before ungraceful restarts and failovers. When bitmap persistence is enabled, non-full QR Volume Creation or Image backup jobs remain non-full after both graceful and ungraceful restart or failover. If bitmap persistence is *not* enabled, non-full jobs remain non-full only after a graceful restart or failover; in the event of an ungraceful restart or failover, the system will force a full backup or QR Volume creation.

The following table summarizes the expected behavior of the Agents after graceful (planned) and ungraceful ("blue screen," crash, etc.) restarts and failovers, depending on how you configure the software:

	Non-Cluster	Cluster
Default	Allows non-fulls* after graceful restarts only. Ungraceful restarts will force a full backup or QR Volume creation.	Allows non-fulls* after graceful restarts or failovers only. Ungraceful restarts or failovers will force a full backup or QR Volume creation.
Configured for Persistence	Allows non-fulls* after both graceful and ungraceful restarts.	Allows non-fulls* after both graceful and ungraceful restarts and failovers.

*Where non-full refers to QR Incremental Updates or Incremental backups.

ENABLE PERSISTENCE

By default, Persistence is enabled on all volumes. If required, you can also enable Persistence on volumes with the Qsnp2Config tool, which is located in `<Install Directory>\Base` on the client and is installed with QSnap® software. For step-by-step instructions, see one of the following:

- Enable Persistence on a Volume
- Enable Persistence and Configure QSnap on a Cluster

DISABLE PERSISTENCE

Persistence is disabled with the Qsnp2Config. For step-by-step instructions, see one of the following:

- Disable Persistence on Volumes
- Disable Persistence on a Cluster

BITMAP LOCATION

The bitmap location is specified during the installation of QSnap functionality. Only NTFS volumes are supported for the bitmap location. For cluster installation, the default location for storing bitmaps is the corresponding shared volume but in case of an upgrade the shared location selected prior to upgrade will be the default location for the bitmaps. You can change the bitmap location using Qsnp2Config. See Change the QSnap Bitmap Location for step-by-step instructions.

ENABLE SAN ENVIRONMENT

Enabling the SAN Environment allows you to continue incremental operations after a device (SAN-attached disk) has been reconnected. This option is enabled by default. Disabling feature this will force a full QR Volume creation in cases where a disconnected device has been reconnected. See Enable SAN Environment for step-by-step instructions.

USE BITMAPS TO MEASURE CHANGE

You can predict the size of your next incremental job by using the TrackBlockIO tool. This tool looks at the bitmap information to determine the amount of changed blocks. See Use TrackBlockIO to measure changed blocks for step-by-step instructions.

QSNAP® DRIVER ON UNIX

The Unix QSnap driver interfaces between File System (UFS) and *sd* (disk driver) on Unix. The Unix QSnap driver preserves all the properties of the underlying *sd* driver. For both applications and file systems, it is transparent. The driver intercepts I/O and keeps track of the blocks that have been modified. The bitmap file location is managed by the Unix QSnap® driver.

- The Unix QSnap driver creates device nodes in `/devices`. It does not have hot pluggable support for devices.
- On Unix, QSnap bitmaps are referred to as CXBF bitmaps, and they are managed by the QSnap Unix driver.
- The Unix QSnap driver comes with a utility called CVSnap, which can be used to configure and list CXBF devices.
- The Solaris EFI disk label, introduced in Solaris 10 to support disk labels above one terabyte, is not supported by the Unix QSnap driver.

CXBF DEVICES

On a Unix platform, a CXBF device is a volume or partition that is monitored by the CXBF block-filter driver. Any volume that you add to a subclient is automatically configured as a CXBF device. You can also create CXBF devices using the **Configure cxbf device** right-click option in Volume Explorer on any volumes available to the client.

- CXBF should not be configured on devices containing the Operating System software. Typically, this is the device mounted to "/".
- For Quick Recovery Agent adding a volume to the subclient content does not automatically configure it to be a CXBF device.
- Raw volumes are not automatically configured as CXBF devices when added to the subclient content. Use the Volume Explorer to configure raw volumes as CXBF devices.
- CXBF devices can only be configured on quiescent file systems (i.e. umountable or not busy).

FOR AIX:

The CXBF devices for AIX disks must be Logical Volume Manager (LVM) devices with the JFS2 virtual file system. (The JFS virtual file system is not supported.) These devices will be created in the paths of:

`/dev/cxbf/cxbf`*x*`/blk`

where *x* is a number from 0 to 32256.

The devices are selected by their device name in the CommCell® Console (for example, `/dev/cxbf/cxbf24/blk/dept`).

- CXBF devices on AIX do not support online file system expansion or shrink using the `chfs` command. If a volume or file system controlled by a CXBF block-filter driver is extended or shrunk, it will fail. Therefore, you must de-configure the CXBF device before using `chfs`.
- If a CXBF device on AIX was configured using Volume Explorer (where the CXBF device was added to `/etc/filesystems`), be sure to mount the CXBF device before you deconfigure it. If an unmounted CXBF device on AIX is deconfigured using Volume Explorer, the `/etc/filesystems` system file will be out of synchronization with the actual file system configuration.

  If `/etc/filesystems` and the actual file system configuration get out of synchronization, a manual edit of `/etc/filesystems` is required to correct the inconsistencies. Remove the cxbf entry and uncomment the original entry (that is, remove "*Galaxy" from the beginning of each line).

FOR SOLARIS:

The CXBF devices for these disks will be created in the paths of `/dev/cxbf/dsk` and `/dev/cxbf/rdsk`. They are selected by their device name in the CommCell® Console (for example, `/dev/cxbf/dsk/c2t1d1s1`).

The naming convention for CXBF devices is as follows:

`cxtxdxsx`

Where the system assigns integer **`x`** as follows:

- **`cx`** is the controller number
- **`tx`** is the target number
- **`dx`** is the LUN
- **`sx`** is the slice number

QSNAP® SOFTWARE PERFORMANCE ON SOLARIS

If the server system is heavily loaded with I/O, multiple processes, or low memory, you may need to tune the buffer cache parameter in `/etc/system`. If the buffer cache is too low, backup processes may hang in heavily loaded condition.

The `bufhwm` variable, set in the `/etc/system` file, controls the maximum amount of memory allocated to the buffer cache and is specified in KB. The default value of `bufhwm` is 0, which allows up to two percent of system memory to be used. This may need to be increased to 10 percent for a dedicated file server with a relatively small memory system, and can be increased up to 20 percent. On a larger system, the `bufhwm` variable may need to be limited to prevent the system from running out of operating system kernel virtual address space.

The buffer cache is used to cache inode, indirect block, and cylinder group related disk I/O only. The following is an example of a buffer cache (`bufhwm`) setting in the `/etc/system` file that can handle up to 10 MB of cache. This is the highest value to which you should set `bufhwm`:

`set bufhwm=10240` (the unit is KB)

FOR LINUX:

Linux disk devices are named as follows:

`/dev/hdx` for an IDE disk

`/dev/sdx` for a SCSI disk

`/dev/vg00/vol0` for a LVM disk

where *x* is a letter a, b, c, d, etc.

The CXBF devices for these disks will be created in the paths of:

`/dev/cxbf/cxbfx/`

where *x* is a number from 0 to 32256.

They are selected by their device name in the CommCell® Console (for example, `/dev/cxbf/cxbf32/dsk`).

To find the relationship between the Linux CXBF device and original device, go to `/proc/cxbf/cxbfx` and type **`cat hdd_name`**. This will report the original device name.

CXBF DEVICE CONFIGURATION

After installing the software, you must create CXBF devices for the volumes you want to back up/protect. When you add a volume to a subclient content, it is automatically configured as a CXBF device.

See the following procedures for step-by-step instructions:

- Configure a CXBF Device in Volume Explorer
- Test a CXBF Device
- Use a Disk that has been Newly Added to a Client
- Change a Disk Label

    - If you encounter errors while configuring or deconfiguring CXBF devices on the client, go to `/cvd.log` on the client for more detailed information concerning the error.
    - If you apply an update to an AIX, Linux, or Solaris platform where CXBF drivers have been installed and CXBF devices have been configured, you must reboot the system after the update is applied in order to activate the CXBF drivers.
    - For Quick Recovery Agent adding a volume to the subclient content does not automatically configure it to be a CXBF device.
    - Use the Volume Explorer to deconfigure a CXBF device, deleting a subclient does not deconfigure CXBF devices for a client.

For Oracle Volumes:

- Edit the Mapping File for RAW Devices
- Use a Newly Configured CXBF Device (RAW Devices)
- Use a Newly Configured CXBF Device (File System)

Using the CVSnap Tool:

- Display All CXBF Devices
- Get CXBF Device Information
- Defunc and Delete a CXBF Device

Back to Top

RECOVERY POINTS

When used with the Quick Recovery® Agent (QR) on Windows or with ContinuousDataReplicator (CDR) on Unix, an additional feature is available for the QSnap service to create Recovery Points, by creating a snapshot of the file system data on a QR™ Volume or Destination machine. For more information, see Recovery Points.

Note: When using QSnap functionality with CDR on Windows, you cannot create recovery points.

Back to Top

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.

RECOVERY POINTS FEATURE

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

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

Back to Top

QSnap on Windows - How To

Topics | How To | Related Topics

Windows File System *i*DataAgent

- Enable QSnap on a Subclient
- Change the COW Cache Size
- Change the COW Cache Location
- Change the extent size for Backup Applications
- Change the extent size for Replication Applications

SAN and Volume Level Agents on Windows

- Change the COW Cache Size
- Change the COW Cache Location
- Change the QSnap Bitmap Location
- Enable Persistence on a Volume
- Enable Persistence on Cluster Volumes
- Disable Persistence on Volumes
- Disable Persistence on Cluster Volumes
- Disable QSnap Block-filter Driver on Specific Volumes
- Change the Write Inactivity Period (WIP)
- Enable SAN Environment

Additional Tools for Windows

- Qsnp2Config is covered in the configuration procedures for QSnap Windows Agents above.
- QST2 Tool
- TrackBlockIO

ENABLE QSNAP ON A SUBCLIENT

Related Topics:

- Subclients - SAN *i*DataAgents
- Subclients - Windows/Unix/Macintosh File Systems and Microsoft Data Protection Manager *i*DataAgent
- QSnap for the Unix File System *i*DataAgent
- QSnap for the Windows File System *i*DataAgent
- Things to Consider when Creating and Configuring File System Subclients

Before You Begin

- Review QSnap for the Windows File System *i*DataAgent and QSnap for the Unix File System *i*DataAgent.
- For Unix File System *i*DataAgents or Windows File System *i*DataAgents, also review Things to Consider when Creating and Configuring File System Subclients.

Required Capability: See Capabilities and Permitted Actions

To enable QSnap on a subclient:

1. From the CommCell Browser, right-click the subclient that you want to back up, and choose **Properties**.
2. For Unix File System *i*DataAgents or Windows File System *i*DataAgents, from the **General** tab, select the **Use QSnap** checkbox.
3. Click **OK**. Your next backup will use QSnap.

CHANGE THE COW CACHE LOCATION

For Windows platform Agents, use the following procedure to change the location of the COW Cache for all volumes on a client.

- For the Image Level ProxyHost *i*DataAgent, this value is set in the Backup Host's Image Level Agent Properties.

- If you have manually set values for the cache location in the registry, those values will take precedence over any values set in the CommCell Console.

Required Capability: See Capabilities and Permitted Actions

To change the COW Cache location:

1. From the CommCell Console, right-click the Client icon for which you are changing the COW Cache location and then click **Properties**.
2. From the Client Computer Properties (Advanced) tab, enter the volume on which to place the COW Cache.
   - Directory Path Mount points are not supported.
3. Click **OK**. The new cache location will be used for the next backup/snapshot/QR job.

CHANGE THE COW CACHE SIZE

For Windows platform Agents, use the following procedure to change the minimum and/or maximum COW cache sizes.

- For the Image Level ProxyHost iDataAgent, this value is set in the Backup Host's Image Level Agent Properties.
- The effective maximum for COW cache size is 90% of the volume's total capacity. Even if you enter a number of megabytes that represents a higher percentage, a snapshot will terminate when the cache file reaches 90% of the volume's total capacity.

Required Capability: See Capabilities and Permitted Actions

To change the COW Cache size:

1. From the CommCell Console, right-click the Agent icon for which you are adjusting the COW Cache parameters and then click **Properties**.
2. From the Agent Properties (General) tab, enter the desired values for the minimum and/or maximum COW Cache size.
3. Click **OK**. The new cache settings will take effect during your next backup/snapshot/QR job.

CHANGE THE EXTENT SIZE FOR BACKUP APPLICATIONS

Use the following procedure to change the backup extent size.

This procedure will require a reboot.

To change the extent size:

1. From the `\Base` directory, launch the **Qsnp2Config**.
2. Click **Update BF Params**, enter the desired value in **Extent size in Bytes for Backup Applications** and click **OK** to continue.
3. Click **OK**. The new extent size settings will take effect after a reboot.

CHANGE THE EXTENT SIZE FOR REPLICATION APPLICATIONS

Use the following procedure to change the replica extent size.

This procedure will require a reboot.

To change the extent size:

1. From the `\Base` directory, launch the **Qsnp2Config**.
2. Click **Update BF Params**, enter the desired value in **Extent size in Bytes for Replication Applications** and click **OK** to continue.
3. Click **OK**. The new extent size settings will take effect after a reboot.

CHANGE THE QSNAP BITMAP LOCATION

Use the following procedure to change the location of the QSnap bitmaps. Prior to changing the bitmap location, consider the following:

- The bitmap location must be an NTFS volume.
- Directory Path Mount points are not supported.
- For standard configurations the volume you select must be a local drive. Do not select any fiber channel attached devices, external SCSI devices, or other external hard disks.
- For cluster configurations, the volume you select must be a shared volume of the same Cluster Group and all volumes on the Cluster Group must use this shared volume for the bitmap file storage.
- A reboot may be required for the changes to take place.
- In some instances, changing the location of the bitmaps may force a full backup job. This will happen if the Bitmap Init value is set 1, if the Bitmap Init value remains 0, then a full backup will not be forced.

To change the location of Block Filter (BF) Parameters of an individual volume:

1. From the `\Base` directory, launch the **Qsnp2Config**, which will display the DEFAULT Parameters configured during install.
2. Click on the **Configure Block Filter on Volumes** button.

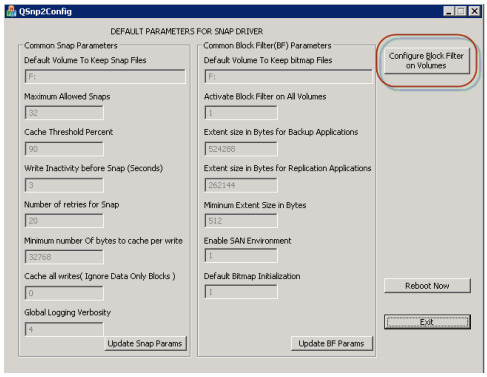

3. In the **Enter a Volume** field, enter the volume you desire to change (for e.g., I:) and click the **Select** button, which will display the current settings for the volume.
4. Click **Deactivate BF**. In the **Bitmap Volume** field, enter the volume where you want to store the bitmap for the selected volume (for e.g., I:).
5. Click the **Activate BF** button.
6. Repeat this process for each volume on which you are going to change the bitmap location.
7. Click **Close** to close the **Configure BlockFilter Volume** dialog box.
8. Click the **Exit** button to close the **Qsnp2Config** dialog box.

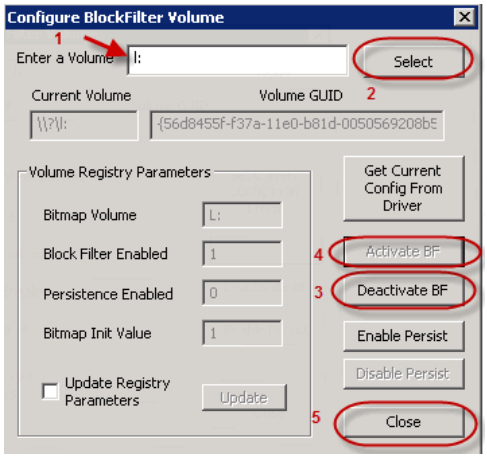

To change the location of Block Filter (BF) Parameters of all volumes:

1. From the `\Base` directory, launch the **Qsnp2Config**, , which will display the DEFAULT Parameters configured during install.
2. Click **Update BF Params**, enter the new bitmap location (for e.g., I:) in the **Default Volume to keep Bitmap Files** field and click **Update**.
3. Click **OK**.
4. Click the **Exit** button to close the **Qsnp2Config** dialog box.

ENABLE PERSISTENCE ON A VOLUME

Use the following procedure to enable persistence on a volume or volumes.

> Disabling persistence does not require a reboot and does not force your next job to be a full.

To enable QSnap bitmap persistence:

1. From the `\Base` directory, launch the **Qsnp2Config**.
2. Click **Configure Volumes**, then enter or **Select** the volume for which you want to enable QSnap bitmap persistence.
3. Click **Enable Persistence**. The Persistence Enabled value should be set to **1**.
4. Repeat this process for each volume on which you are enabling persistence. For the Quick Recovery Agent, remember to configure the destination volumes

in addition to the source volumes.

5. Click the **Exit** button to complete volume configuration.

Enable Persistence and Configure QSnap on Cluster

Once you have installed QSnap, use the following procedure to set persistence and configure QSnap on a cluster.

To enable persistence on a cluster:

1. From the Cluster Administrator software, take the **GxClusPlugin** service resource offline.
2. From the registry editor on both physical nodes, navigate to `HKEY_LOCAL_MACHINE\SOFTWARE\CommVault Systems\Galaxy\Instance<xxx>\iDataAgent`.
   a. Change the DWORD value for `ClusterBFMode` to **2**.
   b. Close the registry editor.
3. Configure your volumes:
   a. From the `\Base` directory, launch **QST2Tool**. This executable can be run from either physical node.
   b. Activate the block-filter driver on the volume.
   c. Enable Persistence.
   d. Set the **Bitmap Init Value** to **0**.
   e. Specify the volume to keep the bitmap. Ensure the volume you specify here is a shared volume.
4. From the Cluster Administrator software, bring the **GxClusPlugin** service resource back online.
5. Initiate a failover by moving the group to the other physical node.
6. The cluster volumes you configured on the original node need to be configured on the current node as well. Once configured on both nodes, you will not need to configure them again. Configure your volumes by repeating Steps 1 - 4 above.

QSnap is now configured for your cluster.

Disable Persistence on Volumes

Use the following procedure to disable persistence on a volume or volumes.

> Disabling persistence does not require a reboot and does not force your next job to be a full.

To disable QSnap bitmap persistence on a volume:

1. From the `\Base` directory, launch the **Qsnp2Config**.
2. Click **Configure Volumes**, then enter or **Select** the volume for which you want to disable QSnap bitmap persistence.
3. Click **Disable Persistence**. The Persistence Enabled value should be set to **0**.
4. Repeat this process for each volume on which you are disabling persistence. For the Quick Recovery Agent, remember to disable persistence the destination volumes in addition to the source volumes.
5. Click the **Exit** button to complete volume configuration.

Disable Persistence on a Cluster

Use the following procedure to disable persistence on a cluster.

To disable persistence on a cluster:

1. From the Cluster Administrator software, take the **GxCVD** and **GxEvMgrc** service resources offline.
2. Add any additional physical disk resources that you want to protect as dependencies to **CVD** (the physical disk resource that QSnap was installed on for the virtual node is listed as a dependency by default).
3. From the registry editor on both physical nodes, navigate to `HKEY_LOCAL_MACHINE\SOFTWARE\CommVault Systems\Galaxy\Instance<xxx>\iDataAgent`.
   a. Change the DWORD value for `ClusterBFMode` to **1**.
   b. Close the registry editor.

4. Configure your volumes:
    a. From the `\Base` directory, launch **Qsnp2Config**. This executable can be run from either physical node.
    b. Click **Configure Volumes**, enter or **Select** the volume for which you want to disable QSnap bitmap persistence.
    c. Click **Deactivate BF** to deactivate the block-filter driver on the volume.
    d. Set the **Persistence Enabled** value to **0**.
    e. Set the **Bitmap Init Value** to **0**.
    f. Set the Bitmap Volume to your desired location for storing the bitmap file. The volume is required to be a local volume.
    g. Click **Reactivate BF** to reactivate the block-filter driver.
    h. Repeat this process for all of the shared volumes you would like to protect on with QSnap on your cluster.
5. From the Cluster Administrator software, bring the GxCVD and GxEvMgrc service resources back online.
6. Initiate a failover by moving the group to the other physical node.
7. The cluster volumes you configured on the original node need to be configured on the current node as well. Once configured on both nodes, you will not need to configure them again. Configure your volumes:
    a. Click **Configure Volumes**, enter or **Select** the volume for which you want to disable QSnap bitmap persistence.
    b. Click **Deactivate BF** to deactivate the block-filter driver on the volume.
    c. Set the **Persistence Enabled** value to **0**.
    d. Set the **Bitmap Init Value** to **0**.
    e. Set the Bitmap Volume to your desired location for storing the bitmap file. The volume is required to be a local volume.
    f. Click **Reactivate BF** to reactivate the block-filter driver.
    g. Repeat this process for all of the shared volumes you would like to protect on with QSnap on your cluster.

Persistence is now disabled on your cluster.

ENABLE/DISABLE QSNAP BLOCK-FILTER DRIVER ON SPECIFIC VOLUMES

Use the following procedure to manually deactivate the QSnap block filter functionality on specific volumes.

To deactivate the block-filter driver on specific volumes:

1. From the `\Base` directory, launch the `Qsnp2Config`.
2. To de-activate the block-filter driver on the volumes you do not want to be monitored:
    a. Click **Configure Volumes**, then enter or select the volume for which you want to disable the block-filter driver.
    b. Click **Deactivate BF**.
    c. Repeat this process for each volume on which you are disabling the block-filter driver.
    d. Click **OK** to continue.
3. Click the **Exit** button to close the `Qsnp2Config`.

CHANGE THE WRITE INACTIVITY PERIOD (WIP)

Use the following procedure to change the write inactivity period. Use caution when changing this value. If the default value of 3 seconds is not sufficient (for example, due to slow disk performance), try small, incremental increases in the value (e.g., increase the value to 5 seconds).

This procedure will require a reboot.

To change the WIP:

1. From the `\Base` directory, launch the **Qsnp2Config**.
2. Click **Update Snap Params**, then enter the new **Write Inactivity before Snap (Seconds)** value and click **OK** to continue.
3. Click the **Exit** button to close the Qsnp2Config.

ENABLE SAN ENVIRONMENT

Use the following procedure to enable SAN environment support, allowing you to perform incremental backups after a SAN device has lost connectivity.

To enable SAN environment support:

1. From the `\Base` directory, launch the **Qsnp2Config**.
2. Click **Update BF Params**, then change the **Enable SAN Environment** value to **1** and click **OK** to continue.
3. Click the **Exit** button to close the Qsnp2Config.

QST2 TOOL

The QST2 tool is used primarily for testing and troubleshooting the QSnap driver. The tool allows you to directly access and use the QSnap driver through command line calls. The QST2 tool is located in the `\Base` directory and has both snapshot and bitmap file commands.

WARNING

- Some features of this tool should not be used without instruction from a support specialist and have been marked **DO NOT USE**.
- This tool provides some options that are better performed using the CommCell Console or Qsnp2Config. If there is another, preferred method to perform a task, it is noted in the description.

To access snapshot commands:

1. From the `\Base` directory, launch the **QST2** tool from the command line by navigating to the `\Base` directory and entering:

   qst2 *x*:

   Where **x** is the drive letter for the volume containing the snapshot(s) with which you intend to work.
2. Enter **1** to access the snapshot options.
3. Select one of the following:
   - **Get Device Snap Info** - measures the amount of data that has been cached on the current snap. Multiply the **FSSectorSize** by the **TotalCachedSectors** to calculate the **Bytes** of data have been cached. This will only return data if a QSnap snapshot already exists on the volume specified when starting the QST2 tool.
   - **Activate Snap** - creates a snapshot of the current specified volume. Note that you must specify a cache partition using options 8 or 13 before attempting to activate.
   - **Deactivate Snap** - deactivates an activated snapshot of the current specified volume.
   - **Expose Snap as Mount Point** - mounts an activated snap to a drive letter.
   - **Remove Exposed Snap Mount Point** - unmounts a snapshot that you had previously mounted.
   - **Test Driver Control Object - DO NOT USE**.
   - **Test IOCTL to QSnap - DO NOT USE**.
   - **Freeze Snap - DO NOT USE.**
   - **Get Cached Extents - DO NOT USE.**
   - **Get Snap Volume Device ID** - displays the Snap Volume Device ID for a snapshot. If you are performing a function that requires the ID, and you do not know the snapshot ID, use this option to get the information.
   - **Add NTFS Cache Partition to QSnap Driver** - specifies an NTFS volume as the preferred cache volume for snapshots.
4. When you are finished working with your snapshots, enter **0** to return to the main menu and then enter **0** again to exit the program.

To access bitmap commands:

1. From the `\Base` directory, launch the **QST2** tool from the command line by navigating to the `\Base` directory and entering:

   qst2 *x*:

   Where **x** is the drive letter for the volume containing the bitmaps with which you intend to work.
2. Enter **2** to access the bitmap options.
3. Select one of the following:
   - **Get info** - displays the information that the QSnap driver has regarding bitmaps for your volume.
   - **Activate Block Filter** - this should be performed from the **Qsnp2Config** and should only be done during specific configuration procedures or for troubleshooting purposes.
   - **Deactivate Block Filter** - this should be performed from the **Qsnp2Config** and should only be done during specific configuration procedures or for troubleshooting purposes.

- **Set Bitmap - DO NOT USE.**
- **Clear Bitmap** - resets the current bitmap information to zero. In a production environment this will cause the next data protection operations to be Fulls. It is not recommended to use this option.
- **Get Bitmap** - prints out the bitmap information to a file. The file will show you which areas of the bitmap are dirty (contain blocks with modified information).

  Select `0` for Image Incremental bitmaps, `1` for Image Differential bitmaps, or `2` for QR Incremental update bitmaps. To write the data to a file, type the file name and path, when prompted to **Press S to dump bitmap on screen**.
- **Get and Merge Bitmap - DO NOT USE.**
- **Enable Persistence** - this should be done from the **Qsnp2Config**. See Enable Persistence on a Volume and Enable Persistence on Cluster Volumes for more information.
- **Disable Persistence** - this should be done from the **Qsnp2Config**. See Disable Persistence on Volumes and Disable Persistence on Cluster Volumes for more information.

4. When you are finished working with your bitmaps, enter `0` to return to the main menu and then enter `0` again to exit the program.

USE TRACKBLOCKIO TO MEASURE CHANGED BLOCKS

Use the following procedures to measure the amount of change to a volume at the block level, over time. Note that previous versions of the TrackBlockIO tool will not work in the current release. Similarly, the TrackBlockIO tool for the current release cannot be used with older versions of the software.

To measure the changed blocks for the next job:

1. From the `\Base` directory, launch the **TrackBlockIO** tool.
2. Enter the drive letter for the volume in the **Enter Drive Letter:** field and click **Make Active Volume**.
3. Select the appropriate bitmap type from the **Bitmap Type** drop-down list:
   - Select **Replica Incremental** to measure block-level changes for QR Agent incremental updates.
   - Select **Backup Incremental** to measure block-level changes for Image Level and Image Level ProxyHost incremental backups.
4. If you want to save the measurements, select the **Save Results To** checkbox and enter a path for the results file (it will use the directory in which the TrackBlockIO is located by default).
5. Click the **Get Changed Data** button to see the point-in-time from which you are measuring change, as well as the amount of change in KB in the **Last Cleared Time** and **Data To Copy** fields, respectively.
6. The value displayed in the **Data To Copy** field is indicative of the size of the next backup or QR Incremental update.

   WARNING:

   Do not enable and use the Clear Bitmap button in a production environment - they are provided for testing purposes outside of production only. See the example below.

If you are **not** performing backups or creating QR Volumes with QSnap, you can use the TrackBlockIO tool in conjunction with the Clear Bitmap option to test the potential size of your jobs before backing up your client. Clearing the bitmaps in a production environment would force a full backup or QR Volume creation.

To measure the changed blocks before backing up a client:

1. Install QSnap, preferably without the Image Level, Image Level ProxyHost or QR Agents.
2. From the `\Base` directory, launch the **TrackBlockIO** tool.
3. Enter the drive letter for the volume in the **Enter Drive Letter:** field and click **Make Active Volume**.
4. Select a bitmap type from the **Bitmap Type** drop-down list. It is not important which type you choose in this scenario, so long as you continue to use the same bitmap type for all of your measurements.
5. Select the **Save Results To** checkbox and enter a path for the results file (it will use the directory in which the TrackBlockIO is located by default).
6. Click the **Get Changed Data** button. A pop-up warning that the Bitmaps are dirty will appear. Click **OK** to continue.
7. Select the **Enable Clear Button** checkbox and click **Clear Bitmaps**.
8. Click the **Get Changed Data** button again. Since the bitmaps have been cleared, your first measured result should be very small.
9. Wait whatever time interval over which you are measuring changed blocks and click the **Get Changed Data** button again. This will display the point-in-time from which you are measuring change, as well as the amount of change in KB in the **Last Cleared Time** and **Data To Copy** fields, respectively.

   For example, if you want to measure the change over a 24 hour period, wait 24 hours and click the **Get Changed Data** button.
10. You can measure the change over different time intervals by repeating Steps 7 through 9, so long as you are not using QSnap in conjunction with any of the supported agents (QR Agent, Image Level or Image Level ProxyHost *i*DataAgents) to back up the client.

Back to Top

QSnap on UNIX - How To

Topics | How To | Related Topics

Unix File System *i*DataAgents

- Mount the COW Cache partition
- Configure a CXBF Device in Volume Explorer
- Deconfigure a CXBF Device in Volume Explorer
- Test a CXBF Device
- Use a Disk that has been Newly Added to a Client
- Change a Disk Label
- Enable QSnap on a Subclient

SAN and Volume Level Agents on UNIX

- Mount the COW Cache partition
- Configure a CXBF Device in Volume Explorer
- Edit the Mapping File for RAW Devices (for Oracle Volumes only)
- Deconfigure a CXBF Device in Volume Explorer
- Test a CXBF Device
- Use a Newly Configured CXBF Device (RAW Devices) (for Oracle Volumes only)
- Use a Newly Configured CXBF Device (File System) (for Oracle Volumes only)
- Use a Disk that has been Newly Added to a Client
- Change a Disk Label

CVSnaptool for UNIX

- Display All CXBF Devices
- Display All Snapshots
- Get CXBF Device Information
- Stop and Delete a CXBF Device on AIX
- Stop and Delete a CXBF Device on Linux
- Defunc and Delete a CXBF Device on Solaris

MOUNT THE COW CACHE PARTITION

When installing the CXBF block filter driver, you were asked to specify a mount point for the copy-on-write cache. Use the example below to mount the partition.

To mount the COW Cache partition:

When you have selected an appropriate volume, mount it to the previously specified mount point using a command similar to the following:

```
mount /dev/cxbf/dsk/c1t1d1s1 /<mount_point>
```

Where:

<mount_point> is the mount point you specified during the installation of the CXBF driver.

- You will be unable to unmount this partition during any job that uses the COW Cache. If you need to change the partition that is mounted to the COW Cache mount point, you must wait for the job(s) to complete.
- For QSnap with the QuickRecovery Agent on Solaris, the COW cache location must be a CXBF device.

CONFIGURE A CXBF DEVICE IN VOLUME EXPLORER

Use the following procedure to create CXBF devices using Volume Explorer.

Before You Begin

- If you encounter errors while configuring or deconfiguring CXBF devices on the client, go to `/cvd.log` on the client for more detailed information concerning the error.

Required Capability: See Capabilities and Permitted Actions

To create a CXBF device:

1. From the **Tools** menu in the CommCell Console, select **Control Panel**; or (if applicable), go to the **Subclient Properties (Content)** tab for your agent.
2. Click or double-click **Volume Explorer** and then click **Yes** at the warning prompt.
3. Select the host connected to the volume you want to configure.
4. Right-click the volume and select **Configure cxbf device** which opens the Submit Configure Request dialog box.
5. The volume path and name are displayed. It is strongly recommended you do not change the extent size for the volume.
    - The default value for extent size is 4096. The minimum value is 128; the maximum value is 524288.
    - The value is based on block size; to convert to bytes, multiply by 512bytes. Examples: 128 x 512bytes = 64KB; 524288 x 512bytes = 256MB
    - For Image Level on Linux the value must be power of 2.
    - In general, smaller values are better for small files and larger values for large files.
6. **Option available only on Solaris**: Enable persistence, if desired.
7. If you want the CXBF device to be automatically mounted after rebooting, you can select the **Add to VFS tab** option (you are asked to select **Unmount Volume before configuration**). This option is only selectable if you are configuring a volume that is mounted to regular device.
8. Click **Ok** to create the CXBF device.

EDIT THE MAPPING FILE FOR RAW DEVICES

Oracle uses an environment variable called DBCA_RAW_CONFIG, and mapping file that tells the Oracle database application where to create data files, redo files and control files. DBCA_RAW_CONFIG is the variable indicating path of mapping file.

The mapping file must be modified to point to the CXBF devices created for the RAW devices.

To edit the mapping file:

1. Open the mapping file.
2. Replace the RAW devices with the appropriate CXBF devices. For example:

    ```
    System=/dev/rdsk/c1t0d2s1
    ```

    Change to:

    ```
    System=/dev/cxbf/rdsk/c1t0d2s1
    ```
3. Ensure the permissions for the devices are set to the Oracle user and Oracle group:

    ```
    chown Oracle:dba /dev/cxbf/rdsk/c1t1d1s1

    chmod 600 /dev/cxbf/rdsk/c1t0d2s1
    ```
4. Start the Oracle database.

DECONFIGURE A CXBF DEVICE IN VOLUME EXPLORER

Use the following procedure to deconfigure CXBF devices using Volume Explorer.

Before You Begin

- If a CXBF device on AIX was configured using Volume Explorer (where the CXBF device was added to `/etc/filesystems`), be sure to mount the CXBF device before you deconfigure it. If an unmounted CXBF device on AIX is deconfigured using Volume Explorer, the `/etc/filesystems` system file will be out of synchronization with the actual file system configuration.

  If `/etc/filesystems` and the actual file system configuration get out of synchronization, a manual edit of `/etc/filesystems` is required to correct the inconsistencies. Remove the cxbf entry and uncomment the original entry (that is, remove "*Galaxy" from the beginning of each line).
- If you encounter errors while configuring or deconfiguring CXBF devices on the client, go to `/cvd.log` on the client for more detailed information concerning the error.

Required Capability: See Capabilities and Permitted Actions

To deconfigure a CXBF device:

1. From the **Tools** menu in the CommCell Console, select **Control Panel**; or (if applicable), go to the **Subclient Properties (Content)** tab for your agent.

2. Click or double-click **Volume Explorer** and then click **Yes** at the warning prompt.
3. Select the host connected to the volume you want to configure.
4. Right-click the volume and select **Deconfigure cxbf device**.
5. Click **Ok** to deconfigure the cxbf device.

ENABLE QSNAP ON A SUBCLIENT

Related Topics:

- Subclients - SAN *i*DataAgents
- Subclients - Windows/Unix/Macintosh File Systems and Microsoft Data Protection Manager *i*DataAgent
- QSnap for the Unix File System *i*DataAgent
- QSnap for the Windows File System *i*DataAgent
- Things to Consider when Creating and Configuring File System Subclients

Before You Begin

- Review QSnap for the Windows File System *i*DataAgent and QSnap for the Unix File System *i*DataAgent.
- For Unix File System *i*DataAgents or Windows File System *i*DataAgents, also review Things to Consider when Creating and Configuring File System Subclients.

Required Capability: See Capabilities and Permitted Actions

To enable QSnap on a subclient:

1. From the CommCell Browser, right-click the subclient that you want to back up, and choose **Properties**.
2. For Unix File System *i*DataAgents or Windows File System *i*DataAgents, from the **General** tab, select the **Use QSnap** checkbox.
3. Click **OK**. Your next backup will use QSnap.

TEST A CXBF DEVICE

Use this procedure to test a CXBF device to see if the driver has successfully attached.

To test the CXBF device:

1. Log on to the client computer as **`root`**.
2. Depending on your operating system, type one of the following commands:
   - For AIX: **`lsdev | grep cxbf`**
   - For Linux: **`lsmod | grep cxbf`**
   - For Solaris: **`modinfo | grep cxbf`**
3. If the CXBF driver was attached successfully the command will return the unique identifier of the driver. For example:
   - For AIX: **`cxbf0 Available N/A`**
   - For Linux: **`cxbf 132200 2`**
   - For Solaris: **`236 781f8000 c13d 234 1 cxbf (cxbf 0.1)`**
4. In the event that the driver did not attach successfully, do the following:
   - For AIX:
     - Execute the **`remdev -dl cxbf0`** command
     - Execute the **`mkdev -c cxbf -s lvdd -t cxbf`** command
     - Run the **`lsdev | grep cxbf`** command again to ensure that the driver is attached
   - For Linux:
     - Execute the **`rmmod cxbf`** command
     - Execute the **`insmod /lib/modules/2.6.18-92.el5/extra/cxbf.ko`** command
     - Run the **`lsmod | grep cxbf`** command again to ensure that the driver is attached
   - For Solaris:
     - Execute the **`rem_drv cxbf`** command
     - Execute the **`add_drv cxbf`** command

- Run the **modinfo | grep cxbf** command again to ensure that the driver is attached

---

USE A NEWLY CONFIGURED CXBF DEVICE (RAW DEVICES)

For the Quick Recovery Agent and Image Level *i*DataAgents, if the Oracle instance has data files residing on a raw partition, soft links need to be re-established after creating the raw CXBF devices.

For example, if an Oracle volume is linked as:

```
data1.dbf -> dev/rdsk/c1t1d1s1
```

To create the links:

1. Shut down the Oracle database.
2. Remove the link:

   rm data1.dbf
3. After installing the CXBF driver, re-establish the link:

   ln -s /dev/cxbf/rdsk/c1t1d1s1
4. Ensure the permissions for the devices are set to the Oracle user and Oracle group:

   chown Oracle:dba /dev/cxbf/rdsk/c1t1d1s1
5. Start the Oracle database.

During Quick Recovery on the QR Agent, the soft links will automatically be associated with the recovery volume. This is analogous to automatically switching the mount points between file system volumes during recovery.

USE A NEWLY CONFIGURED CXBF DEVICE (FILE SYSTEM)

If the selected disk contains an existing file system that is mounted, then the disk will need to be unmounted and mounted on the CXBF device. (All data will be preserved.) For example, if two Oracle volumes were mounted as:

```
/dev/dsk/c1t1d1s1 /ora_data
/dev/dsk/c2t1d1s1 /ora_logs
```

Required Capability: See Capabilities and Permitted Actions

To mount them as CXBF devices:

1. Shut down the Oracle database.
2. Unmount the volumes:

   umount /ora_data
   umount /ora_logs
3. Configure the volumes as CXBF devices using Volume Explorer.
4. After installing the CXBF driver, mount the same volumes as CXBF devices using Volume Explorer:

   mount /dev/cxbf/dsk/c1t1d1s1 /ora_data
   mount /dev/cxbf/dsk/c2t1d1s1 /ora_logs

   After mounting volumes as CXBF devices, do not also mount them as ordinary Solaris devices using the `dev/dsk` path; this might cause a system panic.
5. Give ownership of the devices to Oracle:

   chown Oracle:dba /ora_data
   chown Oracle:dba /dev/cxbf/dsk/c1t1d1s1
   chown Oracle:dba /ora_logs
   chown Oracle:dba /dev/cxbf/dsk/c2t1d1s1
6. Start the Oracle database.

DISPLAY ALL CXBF DEVICES

Use the CVSnap Tool to display all CXBF devices.

To display all CXBF devices:

1. From the command line, type one of the following, depending on your operating system, and press **Enter**:
   - For Solaris: **/usr/sbin/cvsnap**
   - For AIX: **/sbin/cvsnap**
   - For Linux: **/sbin/cvsnap**
2. At the cvsnap> prompt, type **show** (on Solaris) or **show_filters** (on AIX or Linux) and press **Enter**.

   Sample output:

   ```
   No of bf devices 7
   c1t0d26s0
   c1t0d26s1
   c1t0d26s3
   c1t0d26s4
   c1t0d26s5
   c1t0d26s6
   c1t0d26s7
   ```
3. To exit, at the cvsnap> prompt, type **quit** and press **Enter**.

DISPLAY ALL SNAPSHOTS

Use the CVSnap Tool to display all snapshots on the computer.

To display all snapshots:

1. From the command line, type one of the following, depending on your operating system, and press **Enter**:
   - For Solaris: **/usr/sbin/cvsnap**
   - For AIX: **/sbin/cvsnap**
   - For Linux: **/sbin/cvsnap**
2. At the cvsnap> prompt, type **show_snaps** press **Enter**.
3. To exit, at the cvsnap> prompt, type **quit** and press **Enter**.

GET CXBF DEVICE INFORMATION

Use the CVSnap Tool to obtain CXBF device information.

To get CXBF device information:

1. From the command line, type one of the following, depending on your operating system, and press **Enter**:
   - For Solaris: **/usr/sbin/cvsnap**
   - For AIX: **/sbin/cvsnap**
   - For Linux: **/sbin/cvsnap**
2. At the cvsnap> prompt, type one of the following, depending on your operating system, and press **Enter**:
   - **getinfo device=<*deviceID*>**
   - **getinfo minor=<*minor number for device*>**
   - **getinfo bfminor=<*bfminor number for device*>**
3. To exit, at the cvsnap> prompt, type **quit** and press **Enter**.

STOP AND DELETE A CXBF DEVICE ON AIX

Use the CVSnap Tool to stop and delete a CXBF filter.

> After you stop and delete a device, Volume Explorer will no longer be able to detect it; thus the following commands should be run with extreme caution and only when required.

To stop and delete a CXBF device:

1. From the command line, type the following command and press **Enter**.

   /sbin/cvsnap

2. At the **cvsnap>** prompt, type **show_filters** and press **Enter**.

   All the available CXBF filters will be displayed and from this list you can get the filter number which is the number appended with CXBF.

   For example:

   cvsnap>show_filters

   Sample output:

   /dev/cxbf/cxbf1/blk
   /dev/cxbf/cxbf3/blk
   /dev/cxbf/cxbf4/blk

   Here, there are three filters with filter numbers 1,3, and 4.

3. At the **cvsnap>** prompt, type **stop_filter minor=<*filter_minor_number*>** and press **Enter**.

   For example:

   stop_filter minor=1

   Sample output:

   Stopping CXBF node over:
   minor=1

4. At the **cvsnap>** prompt, type **delete_filter_minor=0** and press **Enter**.
5. To exit, at the **cvsnap>** prompt, type **quit** and press **Enter**.
6. Next, the **cxbf-bootscript1.rc** and **cxbf-bootscript2.rc** files in **/sbin** need to be edited. These files are based on CXBF devices present, and are automatically generated every time a new CXBF device is detected. Edit these two files to remove the entries corresponding to the device you just deleted.
7. Verify that the device you just deleted no longer appears, using the procedure detailed in Display All CXBF Devices.

STOP AND DELETE A CXBF DEVICE ON LINUX

Use the CVSnap Tool to stop and delete a CXBF filter.

> After you stop and delete a device, Volume Explorer will no longer be able to detect it; thus the following commands should be run with extreme caution and only when required.

To stop and delete a CXBF device:

1. From the command line, type the following command and press **Enter**.

   /sbin/cvsnap

2. At the **cvsnap>** prompt, type **show_filters** and press **Enter**.

   All the available CXBF filters will be displayed and from this list you can get the filter number which is the number appended with CXBF.

   For example:

   cvsnap>show_filters

   Sample output:

   /dev/cxbf/cxbf0
   /dev/cxbf/cxbf1

   Here, there are two filters with filter numbers 0 and 1.

3. At the **cvsnap>** prompt, type **stop bfminor=<*filter_number*>** and press **Enter**.

   For example:

   stop bfminor=0

   Sample output:

   Stopping CXBF node:
   minor=0

4. At the **cvsnap>** prompt, type **delete bfminor=0** and press **Enter**.
5. To exit, at the **cvsnap>** prompt, type **quit** and press **Enter**.
6. Next, the **cxbf-bootscript1.rc** and **cxbf-bootscript2.rc** files in **/sbin** need to be edited. These files are based on CXBF devices present, and are automatically generated every time a new CXBF device is detected. Edit these two files to remove the entries corresponding to the device you just deleted.
7. Verify that the device you just deleted no longer appears, using the procedure detailed in Display All CXBF Devices.

DEFUNC AND DELETE A CXBF DEVICE ON SOLARIS

Use the CVSnap Tool to defunc and delete a CXBF device.

> After you defunc and delete a device, Volume Explorer will no longer be able to detect it; thus the following commands should be run with extreme caution and only when required.

To defunc and delete a CXBF device:

1. From the command line, type the following command and press **Enter**.

   /usr/sbin/cvsnap

2. At the cvsnap> prompt, type **defunc device=<*deviceID*>** and press **Enter**.

   For example:

   defunc device=c1t0d6s0

   Sample output:

   ```
   defunc device=c1t0d6s0
   device c1t0d6s0 marked defunct.
   ```

3. At the cvsnap> prompt, type **delete device=c1t0d6s0** and press **Enter**.
4. To exit, at the cvsnap> prompt, type **quit** and press **Enter**.
5. Next, the cxbf-bootscript1.rc and cxbf-bootscript2.rc files in /usr/sbin need to be edited. These files are based on CXBF devices present, and are automatically generated every time a new CXBF device is detected. Edit these two files to remove the entries corresponding to the device you just deleted.
6. Verify that the device you just deleted no longer appears, using the procedure detailed in Display All CXBF Devices.

USE A DISK THAT HAS BEEN NEWLY ADDED TO A CLIENT

When new disks are added to a client, and have been labeled, verify that they can be seen. (If not, edit the sd.conf file as shown below.) Then, use Volume Explorer to detect the newly added disks and create cxbf devices on all the volumes of the newly added disk(s).

To edit the sd.conf file:

1. Log on to the client computer as **root**.
2. Navigate to the /kernel/drv/ directory.
3. Edit the file as **vi sd.conf**. The file will need to be modified if the target is beyond 15 and LUN is greater than 7. The following is an example of the sd.conf file:

   ```
   name="sd" class="scsi" target=0 lun=0 to target=15 lun=7
   ```

   You will need to add more entries if required by your hardware configuration.

   The maximum range of target is 0-255 and for each target LUN the range is 0-255.

   > It is strongly recommended that you do not edit the file unless required.

4. See Configure a CXBF Devices in Volume Explorer to configure the volume.

CHANGE A DISK LABEL

When the disk label has changed, for example, after changing sizes of the partitions/slices, the following procedure needs to be followed to see the new CXBF devices with the correct/updated sizes.

Required Capability: See Capabilities and Permitted Actions

To update the CXBF devices:

1. Unmount, deconfigure, and delete the CXBF devices related to that disk, using Volume Explorer.
2. Change the disk label.
3. Detect, configure, and mount the devices using Volume Explorer.

Back to Top

QSnap for the Image Level *i*DataAgent

Topics | Related Topics

OVERVIEW

QSnap can be used with the Image Level *i*DataAgent to ensure that consistent backups can be taken from the active live system with the point-in-time snapshots created. Thus, all of the components necessary for basic Image Level *i*DataAgent functionality are available without requiring specialized hardware.

During an Image Level backup in a QSnap environment, the following sequence of events takes place:

1. The Image Level *i*DataAgent starts a backup operation.
2. On Windows, QSnap snaps the data to be backed up to free space on the source volume or another specified volume. On UNIX, QSnap copies modified data on the cache partition during the snap window. The cache partition can be linked for more information. (The cache partition cannot be on the same volume.)
3. The snapshot is backed up through the LAN.

CONFIGURATION

The following must be done before using the Image Level *i*DataAgent with QSnap:

1. Verify that the computer in which you wish to install the software satisfies System Requirements for the Image Level and File System *i*DataAgents.
2. Install the required software.
3. Review the Hardware and Software Considerations.
4. Setup the `nOverrideVSSSnapSelection` registry key to use QSnap for snapshot creation.

REQUIRED SOFTWARE

For Oracle, the production server must have one of the following installed:

- Oracle 9.2 Database (Enterprise or Standard edition), or
- Oracle 10g Database (Enterprise or Standard edition)

WINDOWS

The following must be installed on the production server:

- Windows File System *i*DataAgent
- Image Level *i*DataAgent
- QSnap

  QSnap is installed automatically during installation of the Image Level *i*DataAgent.

For a **clustered environment**, first install the Windows File System *i*DataAgent and QSnap on the physical nodes, and then reboot the nodes. After rebooting the nodes, install only the Image Level *i*DataAgent on the virtual node. (The Windows File System *i*DataAgent is automatically selected when you install the Image Level *i*DataAgent.)

UNIX

The following software must be installed on the production server:

- Unix File System *i*DataAgent

- Image Level *i*DataAgent
- Unix QSnap driver

For more information and procedures, see Deployment - Image Level *i*DataAgent.

HARDWARE AND SOFTWARE CONSIDERATIONS

WINDOWS

- When working with file systems other than NTFS, you must set an alternate COW Cache location on an NTFS volume. Increase the maximum size of the QSnap COW Cache to accommodate your disk configuration and usage. See Windows COW Cache for more information on COW Cache configuration.
- After installation, it is highly recommended that you configure persistence for the volumes you are protecting. See Persistence for more information.
- Snapshot creation depends on finding, within one minute, a 3-second period of inactivity on the volume being snapped. If the snap is unsuccessful during that one minute period, it must be retried. By default, the number of retry attempts is 3, after which the job is failed. For Image Level on Windows, to increase the chance of finding a suitable write-inactivity period, you can specify a greater number of retry attempts (up to a maximum of 10) by creating and configuring the `nSnapRetry` registry key.

UNIX

- The COW Cache mount point is selected at the time of installation. See Unix COW Cache in the Overview - QSnap page for more information on supported COW cache configurations.
- When you add a volume to a subclient content, it is automatically configured as a CXBF device. See CXBF Devices in the Overview - QSnap page for more information.
    - If you encounter errors while configuring or deconfiguring CXBF devices on the client, go to `/cvd.log` on the client for more detailed information concerning the error.
    - If you apply an update to an AIX, Linux, or Solaris platform where CXBF drivers have been installed and CXBF devices have been configured, you must reboot the system after the update is applied in order to activate the CXBF drivers.
    - Use the Volume Explorer to deconfigure a CXBF device, deleting a subclient does not deconfigure CXBF devices for a client.
    - Raw volumes are not automatically configured as CXBF devices when added to the subclient content. Use the Volume Explorer to configure raw volumes as CXBF devices.
- The Image Level *i*DataAgent for Solaris works with VxVM version 3.1 or higher.

CONFIGURE EXTENT SIZE

Unlike a File System *i*DataAgent, the Image Level Agent backs up extents on the source drive. The extent size is initially set to 512KB for Windows. The extent size can be configured to be larger than 4GB using the BackupExtSize and BackupExtSizeHigh registry keys. For Unix, the default extent size is 2MB (the Unix GUI will display 4096 because 4096 x 512 bytes equals 2MB). In most cases, the default extent size effectively divides the source volume and is best for performance.

But there may be reasons for increasing or decreasing extent size to improve performance. Keep in mind, for example, that a 512KB extent with just 10KB of data is backed up entirely, including its empty blocks of data. Again, this will work fine in most cases, but factors such as the state of data fragmentation on the source, network bandwidth, and server speed should be considered. These factors, along with extent size, can impact both the speed and the size of backups. These factors, along with extent size, can impact both the speed and the size of backups. If for some reasons the extent size is required to be changed, see Change extent size for Backup Applications for more information on changing extent sizes.

If performance is inhibited because of extent size issues, contact your software provider for more information about tuning your software for maximum performance.

Also consider the following implications of changing extent sizes:

- Backup extent size must be uniform across source and destination computers.
- In a clustered environment, all the physical nodes must have the same Backup extent size.
- Backup extent size should not be larger than the size of the volume.
- For Image Level on AIX, Linux or Solaris, the extent size can only be changed when the CXBF device is configured. Thus, to modify the extent size of an existing CXBF device, you must first deconfigure the CXBF device and then configure again.
- When configuring a CXBF device from Volume Explorer for Image Level on Linux, the value entered must be a power of 2. Using a value which is not a power of 2 will cause the CXBF device configuration to fail. By default, the value is 4096; examples of acceptable values would be 128, 256, 512, 1024, 2048, 4096, 8192, 16384, 32768, etc.

Back to Top

VSS for the Image Level *i*DataAgent

Topics | Related Topics

OVERVIEW

The Image Level *i*DataAgent is integrated with the VSS Enabler to conduct backup operations using a VSS software provider. Image Level uses the snapshot to create the back up.

VSS SOFTWARE PROVIDERS

Microsoft Volume Shadow Service is the default, and only supported VSS software provider. The VSS software provider takes a snapshot of the source volume, database or file, which is then used to create backups or QR Volumes (when used with the Quick Recovery Agent).

CONFIGURATION

The following must be done before using VSS.

1. Verify that the production server and Backup Host meet the System Requirements.
2. Install the VSS Enabler.

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.

ENABLER LICENSES

VSS Hardware snapshot also requires a license for the VSS Enabler.

Back to Top

Volume Explorer

Topics | Related Topics

OVERVIEW

Volume Explorer maintains a database of the disk volumes attached to each MediaAgent and Image Level on Unix *i*DataAgents, and the Quick Recovery Agent client, including unique identifying information (WWN/LUN) for each volume. Also, two types of NetApp data server volumes will be displayed in Volume Explorer. Volumes on the primary data server will be displayed under the QR agent, if they are mounted on that client; the volumes on the secondary data server will be displayed under the data server name. Note that the NetApp NAS *i*DataAgent must be installed in order to detect the volumes.

This information can be used to determine when a single physical volume is mapped to different drive letters or paths on multiple hosts. The system can also look up this information during backup and restore operations, eliminating the need for tedious manual entry of WWNs and LUNs for each subclient or restore job.

For Windows clients, only hosts with the QR Agent and MediaAgent installed will appear in Volume Explorer.

For Unix clients, only hosts with the QR Agent and Image Level on Unix *i*DataAgents and the CXBF driver installed will appear in Volume Explorer. For AIX, Solaris and Linux, Volume Explorer will only detect volumes on disks that have been partitioned, or have had a partition table created on them. Data servers with secondary SnapVault licenses will be displayed for use with SnapVault.

See also DisableWWNDetect, a registry key that allows you to change the default WWN detection behavior of the Quick Recovery Agent.

ACCESSING VOLUME EXPLORER

To access Volume Explorer:

1. From the **Tools** menu in the CommCell Console, select **Control Panel**; or (if applicable), go to the **Subclient Properties (Content)** tab for your agent.
2. Click or double-click **Volume Explorer** and then click **Yes** at the warning prompt.

WHEN TO USE VOLUME EXPLORER

Run a Volume Explorer Detect operation for any of the following conditions:

- after installing a new Image Level *i*DataAgent on Unix or QR Agent, but prior to configuring subclients. Valid Volume Explorer volume information is required prior to proceeding with subclient configuration.
- after disk volumes on a Image Level on Unix, or the Quick Recovery Agent client, or a MediaAgent are added, removed, or reconfigured in any way (including partitioning and drive letter assignments). Note a special case for Linux clients: partitions that have been deleted on a locally connected IDE drive will continue to appear in Volume Explorer until the Linux client has been rebooted.
- when configuring NetApp volumes for use with SnapVault.
- for Image Level on Unix, and Quick Recovery Agent on Unix, to configure/deconfigure and mount/unmount CXBF devices.

VOLUME EXPLORER DISPLAY

The Volume Explorer interface consists of two panes. The left pane displays a list of eligible hosts (including all Windows MediaAgents and all clients with Quick Recovery Agent, or Image Level on Unix installed). The right pane lists the disk volumes that Volume Explorer has detected on the currently selected host.

At the top of the window is a toolbar with several available functions; to read more about each of them, refer to the Volume Explorer Help file.

The Detect operation scans for disk volumes currently attached to the selected host and compares their characteristics with the information in the Volume Explorer database. Volume Explorer will automatically add new volumes to the database. If previously configured volumes are not found during the scan, Volume Explorer will ask for permission before removing their records from the database. If the characteristics of a discovered volume do not match those in

the database, Volume Explorer will ask you whether to preserve or update the database records.

Some types of disk hardware do not support the method of WWN discovery implemented by Volume Explorer. In such cases, Volume Explorer may report a Unique ID Type of ATAPI/IDE instead of Fibre Channel WWN. If you want to use a hardware device for SAN-based data movement, you will need to set the appropriate ID type and enter the WWN and LUN for each volume manually. On subsequent detections of the same volume, you should direct Volume Explorer not to update the database, or you will have to repeat this step.

VOLUME EXPLORER VOLUME PROPERTIES

Volume Explorer displays the following properties for each volume:

Device Name	the Windows device name, or Unix mount path for the volume.
Hostname	the hostname of the attached computer.
Scratch Pool	the Scratch Volume Pool to which this disk belongs, if any.
Assign button	allows you to assign the volume to a new or different Scratch Pool.
ID Type	the type of Unique Identifier to be used with this volume. If the disk is SAN-attached, the Fibre Channel WWN type should be used. If it is an internal disk, ATAPI/IDE is appropriate. If it is a virtual disk device created by QSnap, the QSnap ID type will appear here.
ID	a string (based on ID Type) that uniquely identifies the physical disk containing this volume. Note that multiple partitions on the same disk will have the same ID value, but different LB Offset values (see below).
LUN	A Logical Unit Number for the physical disk. This is used in conjunction with the FC WWN ID type to uniquely identify the logical unit on multi-LUN disk devices (e.g. RAID controllers).

Together, the ID Type/ID/LUN represent the unique identity of the physical volume. Some fields (e.g., WWN) are not always detectable by Volume Explorer. Therefore, Volume Explorer allows you to enter the correct values manually, if needed.

LB Offset	the Logical Block Offset at the beginning of the disk partition.
Block Count	the size of the partition in physical disk blocks.
Block Size	the size of a physical disk block in bytes.

If a Quick Recovery Agent volume creation/recovery operation fails, verify that the ID Type, ID, and LUN entries in Volume Explorer are correct.

SPECIAL NOTES ON VOLUME EXPLORER

- Snapshots can be deleted in Volume Explorer. Use caution when deleting snapshots in Volume Explorer. Volume Explorer allows you to delete snapshots from the CommServe database, even if they are being used, or scheduled to be used by a job, causing the job to fail.
- Volume Explorer will display all disk volumes found on the selected host, including those which are not connected to a SAN. Usually these non-SAN volumes will show up as ATAPI/IDE ID devices and LUN as zero.
- Some SAN attached disk devices (e.g. EMC Symmetrix) do not support the WWN discovery method used by Volume Explorer. For these disks, you will need to enter the ID Type, ID, and LUN information manually. You will only need to do this once per volume.
- In the case of a SCSI disk device that is connected to the SAN using a FC-SCSI router or gateway, Volume Explorer may be unable to detect the WWN correctly. Enter the WWN of the gateway, and be sure to use the LUN that is assigned by the gateway.
- In the unusual case where a native FC device is connected to the SAN via a FC-SCSI router or gateway device, be sure to use the WWN and LUN assigned by the gateway device. The Volume Explorer may detect the WWN of the device itself rather than the gateway, so check these entries carefully to make sure that you provide the WWN and LUN assigned by the gateway. (This caveat does not apply to ordinary FC-AL hubs or fabric switches.)
- For Unix volumes, Volume Explorer will only detect volumes on disks that have been partitioned, or have had a partition table created on them.

Back to Top

Backup Job History

Topics | How To | Related Topics

OVERVIEW

You can view the backup and restore history of *i*DataAgents, BackupSets/Instances, and subclients.

The **Backup Job History Filter** dialog box allows you view detailed, historical information about backup jobs. Once you have chosen your filter options, they are displayed in the **Backup Job History** window.

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

From this window, you can right-click a backup job to:

- Browse the data backed up by the backup set or instance from the **Backup Job History** window. This is provided as right-click option for each job. (This menu option, when selected, initiates the **Browse Options** dialog box preset with the values needed to browse the data.)
- Browse the snapshots created during SnapProtect backup
- View items that failed during the backup job
- View details of the backup job
- View files that were not indexed during a backup job that performed content indexing
- View associated media
- View events of the backup job
- View a list of items that were backed up
- View a list of items that were moved to media for a SnapProtect backup job
- View the log files of the backup job.
- View the RMAN log of an Oracle backup job.
- View the BRTools log of a SAP for Oracle job. You can view the BRTools log for only those jobs that were initiated from the CommCell Console.

ITEMS THAT WERE BACKED UP

The **View backup file list** option allows you to view a list of the files that were backed up during a backup job, along with the data sizes of each backed up file. The **View backed up messages** option allows you to view a list of messages that were backed up by using, along with the alias name, display name, email address, sender name, and recipient of each message.

From these windows you can conduct searches based on a particular string, allowing to find particular files quickly and easily.

It is not recommended that this option is used to view a very large list of items that were backed up (such as lists that total over 100,000 items). It is suggested that the Browse option is used to find a list of backed up items in such cases.

See View the Items That Were Protected During a Data Protection Operation for step-by-step instructions.

ITEMS THAT FAILED

The items that failed for a data protection operation include individual files that may fail the job even though a particular job completed successfully. You can determine the degree of success for these jobs using this window.

Filters can be used in conjunction with the "Items That Failed" list on the data protection Job History Report to eliminate backup or archive failures by excluding items which consistently fail that are not integral to the operation of the system or applications. Some items fail because they are locked by the operating system or application and cannot be opened at the time of the data protection operation. This often occurs with certain system-related files and database

application files.

Also, keep in mind that you will need to run a full backup after adding failed files to the filter in order to remove them.

A listing of files and folders that failed is not available for the Quick Recovery Agent, or the Image Level and Image Level ProxyHost *i*DataAgents. These agents do not perform a file level backup/copy.

Certain application related files can never be backed up by the File System *i*DataAgent due to the nature of the data. For example, Microsoft SQL Server database files cannot be backed up by the File System *i*DataAgent. In this and other similar circumstances, consider entering files such as these as exclusions in the corresponding subclient filter.

See View the Items That Failed For a Data Protection Operation for step-by-step instructions.

PRUNING BACKUP HISTORY INFORMATION

You can prune backup history information based on the number of days established in the **Days to keep the backup job histories** option from the **Media Management Configuration (Service Configuration)** dialog box available in the **Control Panel**.

If you have installed the SQL Server *i*DataAgent, do not use the stored procedure **sp_delete_backuphistory, sp_delete_database_backuphistory** and **sp_delete_backup_and_restore_history** provided by Microsoft clean up backup history. By default backup history is automatically pruned from the CommServe database and the Microsoft SQL Server, as necessary.

SUPPORTED FEATURES

- NAS *i*DataAgents do not support the ability to view items that failed.
- The Image Level and Image Level ProxyHost *i*DataAgents do not support the ability to Browse the data of a selected backup job in Backup Job History.

CONTENT INDEXING HISTORY INFORMATION

Content Indexing history can also be viewed of *i*DataAgents, BackupSets/Instances, and subclients. The following information is displayed:

ITEMS THAT WERE SUCCESSFULLY CONTENT INDEXED

You can view the list of items that were successfully content indexed during a Content Indexing operation for a particular job. for step-by-step instructions, see View the Items that Were Successfully Content Indexed.

CONTENT INDEXING FAILURES

Content Indexing failures allows you to look at the messages, files and documents that could not be indexed during a content indexing operation. Content Indexing looks at each file (of the supported data types) and indexes its contents allowing advanced searches of backed up/archived/migrated data.

Files that were not indexed, (perhaps because the file's content could not be read) are added to the Content Indexing Failures list, and are viewable from the View Content Index (Failed Items) option in the Job History window. For step-by-step instruction, see View the Items that Failed to Content Index.

Back to Top

Backup Job History - How To

Topics | How To | Related Topics

VIEW BACKUP JOB HISTORY

To view backup history:

1. From the CommCell Browser, right-click the entity (client computer, *i*DataAgent, backup set or subclient) whose backup history you want to view, click **View**, and then click **View Backup History.**
2. From the Backup History filter window select the filter options, if any, that you want to apply, and then click OK. The system displays the Backup Job History window.
3. Click **OK**.

VIEW THE ITEMS THAT WERE PROTECTED DURING A DATA PROTECTION OPERATION

This option is available for File System-like agents.

Required Capability: none required

To view the list of items that were protected during 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 the necessary options 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 protected items you want to view, and then select **View backup file list/View Backed Up Messages.** The **Backup file List** window displays a list of the backed up files/messages that were included in the backup job. You can use the **Search** option to find items in the window.
4. Click **File** -> **Exit**.
5. Click **Close** from the **Job History** window.

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 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 MEDIA OR MOUNT PATHS OF A JOB HISTORY

To view media or mount paths associated with 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 select the appropriate history.
2. From the Job History window select the filter options, if any, that you want to apply, and then click **OK**.
3. From the job history widow, right-click the backup whose media or mount paths you want to view, and then click **View Media**.
4. The Media Used By Job ID window displays a list of media or mount paths used by the operation.
5. Click **OK**.

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 ITEMS THAT WERE MOVED TO MEDIA DURING SNAPPROTECT BACKUP


This option is available for the SnapProtect Backup.

To view the list of items that were moved to tape during SnapProtect Backup.

1. From the CommCell Browser, right-click the entity whose history of data protection operations you want to view, click **View**, and then click the necessary options 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 items moved to media you want to view, and then select **View Backup Copy file listing.** The **Backup file List** window displays a list of the backed up files that were included in the backup copy job. You can use the **Search** option to find items in the window.

- To view the files moved to media for a backup copy job, right-click the SnapProtect backup job corresponding to the Backup Copy job and select **View Backup Copy file listing**.
- View backup items will not display anything for a Backup Copy job.

4. Click **File** -> **Exit**.
5. Click **Close** from the **Job History** window.

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.

VIEW THE ITEMS THAT WERE SUCCESSFULLY CONTENT INDEXED

This option is available for operations that performed content indexing.

To view the list items that were not indexed during content indexing:

1. From the CommCell Browser, right-click the entity whose operations you want to view, click **View**, and then click the necessary options 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 job for which you want to view the successfully content indexed items, select **View Content Index**, and click **Successful Items**.
4. Click **Close**.
5. Click **Close** from the **Job History** window.

VIEW THE ITEMS THAT FAILED TO CONTENT INDEX

This option is available for operations that performed content indexing.

To view the list of items that failed to content index:

1. From the CommCell Browser, right-click the entity whose operations you want to view, click **View**, and then click the necessary options 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 job for which you want to view the list of items failed to content index, select **View Content Index**, and click **Failed Items**.
4. Click **Close**.
5. Click **Close** from the **Job History** window.

RESUBMIT A BACKUP JOB

To resubmit a backup job:

1. From the CommCell Browser, right-click the subclient whose backup history you want to view, click **View**, and then click **View Backup History.**

   Additionally, you can view the backup history for a client computer, *i*DataAgent, or backup set . However, the dialogs displayed may be different.

   Note, if viewing the backup history for a client computer, right-click the comptuer name and select **Job History.**

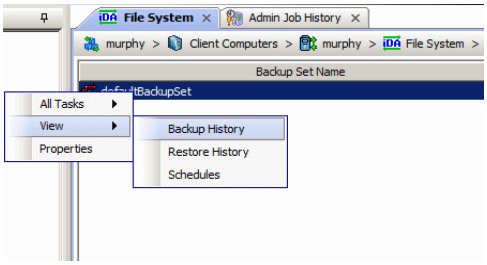

2. From the Backup History filter window select the filter options, if any, that you want to apply, and then click OK. The system displays the Backup Job History window.

   Note: If viewing the job history for a client computer, ensure that the **Backup** radio button is selected.

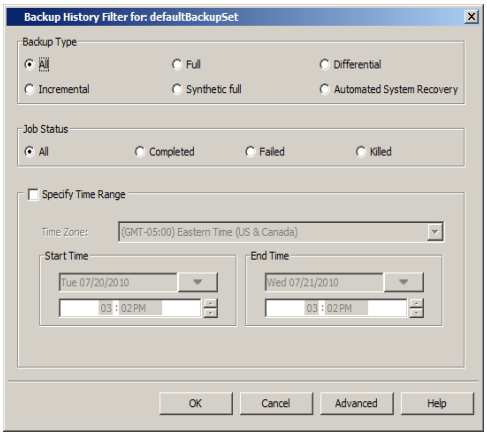

3. The Backup Job History window displays with the specified filter options.

4. Right-click on any job, and select **Resubmit Job**.

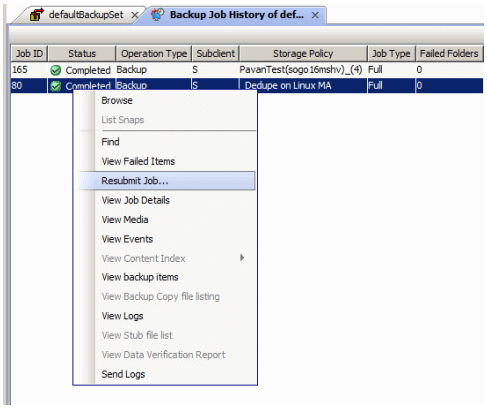

5. From the Backup Options dialog box, select the job options appropriate for the job you want to restart.

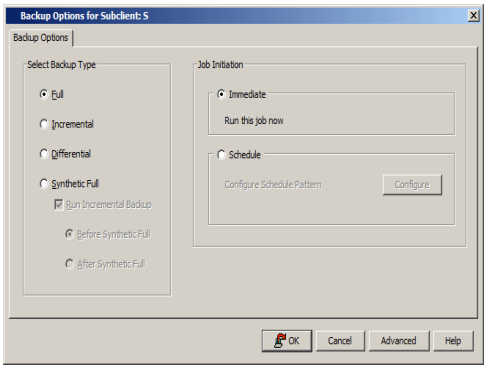

6. If you need to run the backup operation immediately, select **Immediate** from the **Job Initiation** tab. Go to step 11.

7. If you need to schedule the restore operation, select **Schedule** from the Job Initiation tab and click **Configure**.

8. From the **Schedule Details** dialog box that appears, select the appropriate scheduling options.

   Click **Options** to view the Advanced Schedule Options dialog box.

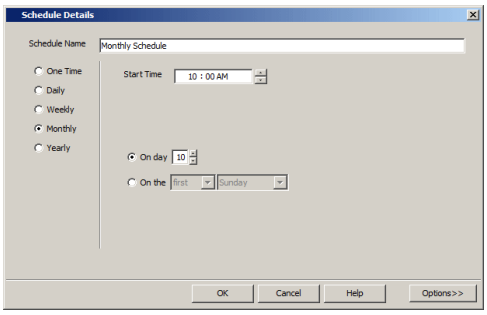

9. From the **Advanced Schedule Options** dialog box:
   - **Range of recurrence**: Specify the date on which you want this schedule to take effect.
   - **Repeat**: Select the value for which you want to run the job repeatedly on the day in which the job is scheduled to run.
   - **Time Zone**: Select a specific time zone from which the job schedule time will be based.

   Click **OK** to close the **Advanced Schedule Options** dialog box**.**

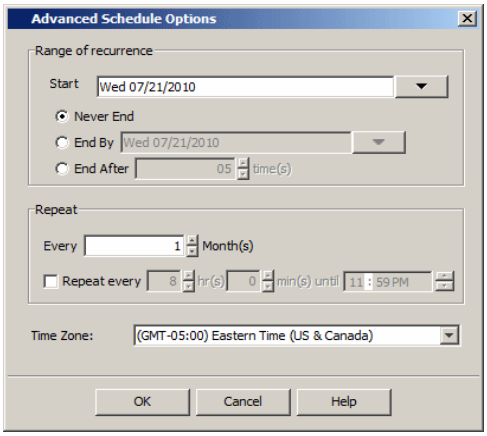

10. Click **OK** to close the **Schedule Details** window.
11. Click **OK** to close the job restart window.

Back to Top

Restore Job History

Topics | How To | Related Topics

OVERVIEW

The **Restore History Filter** dialog box allows you to view detailed, historical information about restore jobs.

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

Once you have chosen your filter options, they are displayed in the **Restore Job History** window. From this window you can right-click a restore job to:

- View Restore Items; items in the job that were **Successful**, **Failed**, **Skipped** or **All**. These items, if any, will be listed in the **Restored Files** window.
- View Job Details of the restore job. The job details will be listed in the **Job Details** window.
- View Events of the restore job. The job events will be listed in the **All Found Events** window.
- View Log files of the restore job. The job log files will be listed in the **Log File** window.
- View the RMAN Log of an Oracle restore job. The RMAN Log will be listed in the **Oracle Restore Log** window.
- View the BRTools log of a SAP for Oracle restore job. You can view the BRTools log for only those jobs that were initiated from the CommCell Console.

ITEMS THAT ARE RESTORED

When viewing files that are restored in the **Restored Files** window, each of the files is listed with the restore status level appended at the end of the file path. The possible status levels are: `RESTORED`, `FAILED` and `OLDER`.

Successfully restored files will be listed with `RESTORED` appended to the file path. If files are not restored/recovered due to errors, the file paths will be appended with `FAILED`. Under some circumstances, the system may not restore/recover certain files because they are older versions of the same files already present in the files system; these files are appended with the word `OLDER`.

SUPPORTED FEATURES

Consider the following.

- NAS *i*DataAgents do not support the ability to view failed/successful item lists.
- Restore Job History will not display Oracle `rman_util` jobs at the instance level.

Back to Top

Restore History - How To

Topics | How To | Related Topics

VIEW RESTORE JOB HISTORY

To view the restored items associated with a job:

1. From the CommCell Browser, right-click the entity (e.g., subclient, client computer, etc.) whose job restore history you want to view, click **View**, and then click **Restore 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 restored items you want to view; click **View Restore Items**, and select from the type of items to view: **Successful**, **Failed**, **Skipped** or **All**.
4. The **Restored Files** window will display the selected type of restored items for the job.
5. Click **OK**.

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 MEDIA OR MOUNT PATHS OF A JOB HISTORY

To view media or mount paths associated with 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 select the appropriate history.
2. From the Job History window select the filter options, if any, that you want to apply, and then click **OK**.
3. From the job history widow, right-click the backup whose media or mount paths you want to view, and then click **View Media**.
4. The Media Used By Job ID window displays a list of media or mount paths used by the operation.
5. Click **OK**.

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