

User Guide - SQL Server *iDataAgent*

TABLE OF CONTENTS

OVERVIEW

Introduction

Key Features

Add-On Components

Terminology

NEW FEATURES

SYSTEM REQUIREMENTS

SUPPORTED FEATURES

GETTING STARTED

SQL SERVER DEPLOYMENT

Interactive Install

Install Software from CommCell Console

SQL SERVER DEPLOYMENT ON A CLUSTER

SQL SERVER IDATAAGENT DEPLOYMENT ON A NON-CLUSTER

Install SQL Server iDataAgent in Veritas Cluster Environment

Install SQL Server iDataAgent in HP Scalable NAS/Polyserve Cluster Environment

SQL SERVER CONFIGURATION

SQL SERVER BACKUP

SQL SERVER RESTORE

ADVANCED

SQL SERVER CONFIGURATION

Understanding the CommCell Console

Creating User-Defined Subclients

Managing Instances

Enabling/Disabling Automatic Database Discovery

Automatically Discovering Databases in Offline States

Excluding Databases from Backups

Managing Databases Deleted from SQL Server

Setting Up Backup Conversion Rules

Enhancing Performance During Backups

Configuring the Number of Log Backups to Run Before a Full Backup

Configuring Log Backups to Run Without Full Backups

Configuring Backups of Locked Databases Using VSS

Configuring Data Streams

Configuring User Accounts for Backups

Modifying an Agent, Instance, or Subclient

Deleting an Agent, Instance, or Subclient

SQL SERVER BACKUP

Full Backups

Transaction Log Backups

Differential Backups

Compressing Backups with Native SQL Compression (for SQL Server 2008 or later)

Perform Partial Backup to Exclude Read-Only File Groups

Scheduling a Backup

Managing Jobs

Additional Options

SQL SERVER BROWSE DATA

Understanding the Browse Window

Browsing Data

Browse from Copies

SQL SERVER RESTORE

Restoring a Database

Recovering a Database

Restoring System Databases

Restoring Files or File Groups

Restoring SQL Server Transaction Logs (Step Restore)

Applying Log Backups

Restoring SQL Databases to a Point in Time

Restoring Only the Logs

Restoring SQL Server Databases to a Transaction Mark

Restoring a Database With Restricted Access

Preserving Replication Settings

Restoring the Captured Changes (For SQL Server 2008 and Above)

Increasing the VDI Timeout

Performing an Out-of-Place Restore

Browsing Data

Listing Media

Scheduling a Restore

Setting Up Pre-Post Processes

Managing Restore Jobs

Additional Restore Options

REPLICATION

Replication Using Using Warm Database Restore

DATA AGING

Getting Started

Extended Retention Rules

Data Aging for Transaction, Archive, and Logical Log Backups

Stored Procedures

SQL Back in Time Restores and Data Aging Rules

Prune SQL Logs Only By Days

DISASTER RECOVERY

Planning for a Disaster Recovery

Rebuilding the Operating System

Rebuilding the SQL Server Instance

Restoring the SQL Server Instance

ADDITIONAL OPERATIONS

Audit Trail

Auxiliary Copy

License Administration

Online Help Links

Operating System and Application Upgrades

Operation Window

Schedule Policy

Storage Policy

Uninstalling Components

BEST PRACTICES

FREQUENTLY ASKED QUESTIONS

BACKUP TROUBLESHOOTING

RESTORE TROUBLESHOOTING

Overview - SQL Server iDataAgent

TABLE OF CONTENTS

Introduction

Key Features

- Full Range of Backup Options
- Advanced SQL Server Restore Capabilities
- Efficient Job Management and Reporting
- Backup and Recovery Failovers
- Block Level Deduplication
- SnapProtect Backup

Terminology

INTRODUCTION

The Microsoft SQL Server iDataAgent provides a simplified end-to-end backup and recovery solution for SQL data in your enterprise. The product can be used to perform both full system rebuilds and granular recovery of the data.

KEY FEATURES

FULL RANGE OF BACKUP OPTIONS

The SQL iDataAgent provides the flexibility to backup the SQL database from different environments. You can perform a full or incremental backup of the entire instance, individual databases or files and file groups, and the transaction logs at any point of time as described below:

DATABASE BACKUPS


You can backup both the system and user-defined databases. You can comprehensively backup all the databases in an instance or schedule backups for the individual databases. You can also auto-discover new databases to comprehensively manage the backup of all databases in your environment.

TRANSACTION LOG BACKUPS

Transaction log backups captures the transaction log whether the transaction was committed or not. The use of transaction log backups make point in time recovery possible. You can restore to any point in time within the transaction log.

FILE AND FILE GROUPS BACKUPS

Files or file group backups allows you to backup individual files or file groups. This functionality can be critically important, especially for large databases. Whereas a full database backup captures all files of a given database, file and file group backups allow you to back up selected portions of a database individually. As with database backups, the system provides the option of performing full, differential, and transaction log backups of file and file groups. Note that when running a transaction log backup for a File/File Group subclient, the database log is automatically backed up.



ADVANCED SQL SERVER RESTORE CAPABILITIES

The SQL iDataAgent provides the ability to recover databases or entire SQL instance. There is no mounting, no recovery wizards, no extra steps needed – the software takes care of it all. This includes the following abilities:

- Full or Partial Restore databases
- Restore and replay transaction logs
- Set Database state during restore (Recovery, Standby, No Recovery)
- Point-in-time recovery

EFFICIENT JOB MANAGEMENT AND REPORTING

You can view and verify the status of SQL backup and recovery operations from the Job Controller and Event Viewer windows within the CommCell Console.

WHERE TO GO NEXT

Install the SQL iDataAgent

Walks you through the process of installing the SQL iDataAgent.

You can also track the status of the jobs using Reports, which can be saved and easily distributed. Reports can be generated for different aspects of data management. You also have the flexibility to customize the reports to display only the required data and save them to any specified location in different formats. For example, you can create a backup job summary report to view at-a-glance the completed backup jobs.

In addition, you can also schedule these reports to be generated and send them on email without user intervention.

BACKUP AND RECOVERY FAILOVERS

In the event that a MediaAgent used for the backup or recovery operation fails, it is automatically resumed on alternate MediaAgents. In such cases, the backup or restore job will not restart from the beginning, but will resume from the point of failure. This is especially useful for backups and restores on large SQL databases.

In the event, that a network goes down, the backup and recovery jobs are resumed on alternate data paths. Similarly, in the event of a device failure, the jobs are automatically switched to alternate disk and tape drives.

BLOCK LEVEL DEDUPLICATION

Deduplication provides a smarter way of storing data by identifying and eliminating the duplicate items in a data protection operation.

Deduplication at the data block level compares blocks of data against each other. If an object (file, database, etc.) contains blocks of data that are identical to each other, then block level deduplication eliminates storing the redundant data and reduces the size of the object in storage. This way dramatically reduces the backup data copies on both the disk and tapes.

SNAPPROTECT BACKUP

SnapProtect Backup enables you to create a point-in-time snapshot by temporarily quiescing the data, taking a snapshot, and then resuming live operations. SnapProtect backups work in conjunction with hardware snapshot engines.

TERMINOLOGY

The SQL documentation uses the following terminology:

CLIENT	The computer in which the iDataAgent is installed and contains the data to be secured.
INSTANCE	The SQL instance used for backup and restore operations.
SUBCLIENT	The SQL databases within the instance used for backup and restore operations.



New Features - SQL Server iDataAgent

NEW FEATURES FOR MICROSOFT SQL SERVER IDATAAGENT

DATA PROTECTION OPERATIONS

- For SQL backups, several new SQL settings are now available. Adjusting these parameters may improve performance and you can enable or disable log consistency checking. You can accept either the defaults or change the values for parameters such as Block Size, Buffer Count, Maximum Transfer Size, and Log Consistency Checking. For more information, see [Enhancing Performance During Backups](#).
- Backups for SQL Server 2008 Enterprise or later can now be compressed by SQL Server before being backed up using Calypso. Backups will be smaller, which should significantly increase the backup speed. However, VSS-enabled and Snap backup types are not supported with this feature. For more information, see [Compressing Backups with Native SQL Compression](#)
- The SQL Server iDataAgent now can restore backups of replicated databases. With this feature, a published database can be restored to a server other than the server where the database was created. This is relevant only to Calypso full backups and if the database was part of a replication implementation. For more information, see [Preserving Replication Settings](#)
- A new backup conversion rule for SQL Server iDataAgent is now available that when configured, will not convert the log backup to a Full backup if the log backup was performed using software other than Calypso. For more information, see [Setting Up Backup Conversion Rules](#)
- You can configure the size of the application data fetched for data transfer during backup operations. Tuning the application read size in alignment with the source application's buffer allocation increases the rate of data transfer during backup. See [Advanced - Microsoft SQL Server Configuration](#) on configuring the application read size for more information.
- When SQL databases were automatically discovered but later deleted from an application (e.g., SQL Server), they will be automatically deleted from the content of the default subclient when the next backup is run. This alleviates having to delete the databases manually. Databases that are manually added and later deleted from an application are not automatically removed from the subclient when the next backup runs and the backup job will complete with errors. See [Managing Databases Deleted from SQL Server](#) for more information.
- By default, SQL databases are automatically discovered and assigned to the default subclient. You can disable this auto-discovery with a new subclient properties option if you need more control of the SQL databases data protection strategy and wish to manually assign them to subclients instead. For more information, see [Enabling/Disabling Automatic Database Discovery](#).

DATA RECOVERY OPERATIONS

SQL databases can now be restored if they were backed up with the Change Data Capture property enabled. This is a useful SQL Server 2008 feature in that all change activity (insert, update, and delete) is captured and applied to SQL tables. The details of the changes are available in an easily consumed relational format. See [Restoring the Captured Changes](#) for more information.

NEW COMMCELL FEATURES SUPPORTED FOR MICROSOFT SQL SERVER IDATAAGENT

DEPLOYMENT

CUSTOM PACKAGE

The Custom Package feature is now extended to almost all products in the Calypso suite. Using Custom Packages, you no longer have to push the entire software DVD through a network, which is especially useful for reducing WAN/LAN payload while installing remote clients.

It is also possible to create Custom Packages using a customized .xml file.

When used in conjunction with the Install Software from the CommCell Console and Automatic Updates features, WAN bandwidth can be drastically reduced during remote site installations. [Learn more...](#)

INSTALL FROM THE COMMCELL CONSOLE

The software installation for this component can be initiated and managed from the CommCell Console, which facilitates the building of your CommCell and eliminates the need to manually install the software. Additionally, the installation of this component can be scheduled to occur at a time suitable for your environment. For more information, see [Install Software from the CommCell Console](#).

UNINSTALL FROM THE COMMCELL CONSOLE

This component can be uninstalled using the CommCell Console. The Uninstall Software utility allows you to quickly see a list of the software packages installed on the selected computer, from which you can then select the components to uninstall. With this, you can easily manage removing software components from client computers and MediaAgents in your CommCell without having to directly access each computer. For more information, see [Uninstall Components using the CommCell Console](#).

ADDITIONAL SNAP ENGINE SUPPORT FOR SNAPPROTECT

The following Snapshot Engines are now supported for SnapProtect:

- Data Replicator
- Dell EqualLogic
- HP StorageWorks EVA
- IBM XIV
- EMC Celerra

SNAP TEST TOOL

Snap Test tool is now available to test basic snap engine operations. See SnapProtect - Snaptest Tool, for more information.

MULTI-STREAMING FOR SNAPPROTECT

Multi-stream backups are now supported for SnapProtect. You can now use multi-stream when moving data to media.

VOLUME MANAGER SUPPORT

SnapProtect Volume Manager support has been extended to support more configurations for e.g., Multiple Physical Volumes containing one Logical Volume. See Supported Volume Managers, for a complete list of volume managers supported for SnapProtect.

[Back to Top](#)

System Requirements - Microsoft SQL Server iDataAgent

System Requirements | **Supported Features**

The following requirements are for the Microsoft SQL Server iDataAgent:

APPLICATION

Microsoft SQL Server 2005 Editions up to the latest Service Pack
Microsoft SQL Server 2008 Editions up to the latest Service Pack
Microsoft SQL Server 2008 R2 Editions up to the latest Service Pack
Microsoft SQL Server 2012 Editions*

*See Considerations for SQL Server 2012.

OPERATING SYSTEM

WINDOWS SERVER 2012

Microsoft Windows Server 2012 Editions

WINDOWS 7

Microsoft Windows 7 Editions

WINDOWS SERVER 2008

Microsoft Windows Server 2008 Editions with a minimum of Service Pack 1*

* Core Editions not supported.

WINDOWS VISTA

Microsoft Windows Vista Editions

WINDOWS SERVER 2003

Microsoft Windows Server 2003 Editions with a minimum of Service Pack 1

WINDOWS XP

Microsoft Windows XP Editions with a minimum of Service Pack 3

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

265 MB of minimum disk space is required for installing the software.

50 MB of additional hard disk space for log file growth.

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

MEMORY

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

PROCESSOR

All Windows-compatible processors supported

PERIPHERALS

DVD-ROM drive

Network Interface Card

MISCELLANEOUS

NETWORK

TCP/IP Services configured on the computer.

NOTES ON MICROSOFT SQL SERVER /DATAAGENT INSTALLATION

- The File System iDataAgent will be automatically installed during installation of this software, if it is not already installed. For System Requirements specific to the File System iDataAgent, refer to System Requirements - Microsoft Windows File System iDataAgent.
- If the option to install SQL Server is disabled make sure there is at least one network protocol enabled. For SQL 2005 go to SQL Server Configuration Manager and enable protocols for each instance.
- Microsoft SQL Server Management Objects (SMO) is required on the client machine. By default, SMO is automatically installed with SQL Server 2005 or later versions.

NOTES ON ALWAYSON PROVIDED BY SQL 2012

SQL Server 2012 comes with a new implementation on clusters named AlwaysOn. If you want to protect the AlwaysOn configuration, consider the following:

- Install the SQL iDataAgent on all physical nodes of the cluster to protect the SQL instances separately. Note that if the active node, which is part of the Availability Group, performs a successful backup then target databases synchronizing with the active node may have errors during backup.
- If you want to restore any databases that are part of the Availability Group, you must remove the database from the Availability Group. Once the restore completes, you can add the database back.

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

Supported Features - Microsoft SQL Server iDataAgent

System Requirements | Supported Features

The following table lists the features that are supported by this Agent.

FEATURE	SUB-FEATURE	SUPPORT	COMMENTS
ADVANCED BACKUP/ARCHIVE OPTIONS	Data tab - Catalog	✓	
	Data tab - Create New Index		
	Data tab - Verify Synthetic Full		
	Job Retry tab	✓	
	Media tab - Allow other Schedule to use Media Set	✓	
	Media tab - Mark Media Full on Success	✓	
	Media tab - Reserve Resources Before Scan		
	Media tab - Start New Media	✓	
	Startup tab	✓	
	VaultTracking tab	✓	
	Comments	✓	Additional options in the Data tab : <ul style="list-style-type: none"> Start a Transaction Log Backup After Successful Backup Perform a Partial Backup to Exclude Read-only File Groups Back up the Tail of a transaction Log Data tab - Catalog options are supported for SnapProtect Backup for this agent.
ADVANCED FILE SYSTEM IDATAAGENT OPTIONS	Automatic File System Multi-Streaming		
	On Demand Data Protection Operation		
	Restore by Jobs		
	Restore Data Using a Map File		
	Comments		
ALERTS AND MONITORING	Global Alerts	✓	
	Job-Based Alerts*	✓	
	Comments		
AUTOMATIC UPDATES	Automatic Updates	✓	
	Comments		
BACKUP/ARCHIVE OPTIONS	Differential Backup	✓	
	Full Backup	✓	
	Incremental Backup		
	Other Backup Types		
	Synthetic Full Backup		
	Comments	✓	Transaction Log Backups Do Not Truncate Log SnapProtect Backup
BACKWARD COMPATIBILITY	Version 7	✓	
	Version 8	✓	
	Version 9		
	Comments		
BROWSE	Browse from Copy Precedence	✓	
	Browse the Latest Data	✓	
	Exclude Data Before	✓	
	Find		
	Full Backup Transparent		

	Browse		
	Image Browse		
	No Image Browse		
	Page Size		
	Specify Browse Path		
	Specify Browse Time	✓	
	Subclient Browse		
	Use MediaAgent		
	View All Versions		
	Comments		
CLUSTERING	Netware cluster		
	Unix Cluster		
	Windows - Microsoft Cluster (MSCS)	✓	
	Windows - Non-Microsoft Cluster	✓	
	Comments	✓	Microsoft Cluster supports Geo-Dispersed Cluster Non-Microsoft Cluster supports VERITAS Cluster and HP Scalable NAS/PolyServe Cluster
COMMAND LINE INTERFACE	Command Line Interface	✓	
	Comments	✓	Qcreate backupset and Qdelete backupset are not supported. An out-of-place restore is supported using Save as Script. For more information, refer to Command Line Interface
COMMAND LINE INTERFACE - SPECIFIC COMMANDS	Qcreate - Backup set/SubClient	✓	
	Qcreate - Instance	✓	
	Qdelete - Backup set/Subclient	✓	
	Qdelete - Client/Agent	✓	
	Qlist globalfilter		
	Qmodify - instance		
	Qoperation - Backup	✓	
	Qoperation - move		
	Qoperation - Restore	✓	
	Comments	✓	Qcreate backupset and Qdelete backupset are not supported. For Qoperation-Restore , an out-of-place restore is supported using Save as Script or Argument File. For more information, refer to Command Line Interface.
COMMCELL MIGRATION	CommCell Migration	✓	
	Comments	✓	CommCell Migration is not supported with SnapProtect backup when using Data Replicator snapshot engine.
CONTENT INDEXING	Offline Content Indexing		
	Comments		
DATA AGING	Basic Retention Rules	✓	
	Extended Retention Rules	✓	
	Unique Data Aging Rules	✓	
	Comments	✓	For the Microsoft SQL Server iDataAgents, data backed up through file/file group subclients cannot be pruned through extended retention rules.
DATA CLASSIFICATION ENABLER	Data Classification Enabler		
	Comments		
DATA COMPRESSION	Client Compression	✓	
	Hardware Compression	✓	
	MediaAgent Compression	✓	
	Comments		
DATA ENCRYPTION	Data Encryption Support	✓	
	Third-party Command Line Encryption Support		
	Comments		
DATA MULTIPLEXING	Multiplexing	✓	
	Comments	✓	For Data Multiplexing , the SQL Server does not support the

			multiplexing of streams "with each other" on to the same media when a data protection operation is performed using the CommCell Console. However, the streams will be multiplexed when the operation is performed using the Command Line Interface.
DEDUPLICATION	MediaAgent Deduplication	✓	
	Source Deduplication	✓	
	Comments		
ERASE BACKUP/ARCHIVED DATA	Erase Data by Browsing		
	Erase Stubs		
	Comments		
GLOBAL FILTERS	Global Filters		
	Comments		
INSTALLATION	Custom Package	✓	
	Decoupled Install	✓	
	Remote Install	✓	
	Restore Only Agents	✓	
	Silent Install	✓	
	Comments	✓	To install this agent as restore only, see Installing Microsoft SQL Server as Restore Only
INSTALLING 32-BIT COMPONENTS ON A MICROSOFT WINDOWS X64 PLATFORM	Install 32-bit On x64	✓	
	Comments		
JOB RESTART - DATA PROTECTION	Not Restartable		
	Restarts from the Beginning		
	Restarts from the Beginning of the Database		
	Restarts from the Point-of-Failure	✓	
	Comments		
JOB RESTART - DATA RECOVERY	Not Restartable	✓	
	Restarts from the Beginning		
	Restarts from the Beginning of the Database		
	Restarts from the Point-of-Failure		
	Comments		
LIST MEDIA	List Media Associated with a Specific Backup Set or Instance	✓	
	List Media Associated with Index		
	List Media Associated with Specific Files and/or Folders		
	List Media Associated with Specific Jobs		
	Comments		
MULTI INSTANCING	Multi Instance		
	Comments		
PRE/POST PROCESSES	Pre/Post Process with Data Protection and Recovery	✓	
	Comments		
RESTORE/RECOVER/RETRIEVE DESTINATIONS	Cross-Application Restores (Different Application version)	✓	
	Cross-Platform Restores - Different Operating System		
	Cross-Platform Restores - Same Operating System - Different Version	✓	
	In-place Restore - Same path/ destination - Same Client	✓	
	Out-of-place Restore - Different path/ destination	✓	
	Out-of-place Restore - Same	✓	

	path/ destination - Different Client		
	Restore Data Using a Map File		
	Restore to Network Drive /NFS-Mounted File System		
	Comments	✓	See Advanced - Microsoft SQL Server Restore.
RESTORE/RECOVER/RETRIEVE OPTIONS	Automatic Detection of Regular Expressions		
	Filter Data From Recover Operations		
	Rename/ Redirect Files on Restore	✓	
	Restore Data Using Wildcard Expressions		
	Restore Data with Pre/Post Processes	✓	
	Restore from Copies	✓	
	Skip Errors and Continue		
	Use Exact Index		
	Use MediaAgent	✓	
	Comments		
RESTORE/RECOVER/RETRIEVE OVERWRITE OPTIONS	Overwrite Files		
	Overwrite if file on media is newer		
	Restore only if target exists		
	Unconditional Overwrite	✓	
	Unconditionally overwrite only if target is a DataArchiver stub		
	Comments	✓	Unconditional Overwrite for existing messages.
SCHEDULE POLICY	Agent Specific Data Protection Schedule Policy	✓	
	All Agent Types Schedule Policy	✓	
	Comments		
STORAGE POLICIES	Incremental Storage Policy*	✓	
	Standard Storage Policies	✓	
	Comments	✓	Incremental Storage Policy does not support Transaction Log backups for this agent. One storage policy can be created for full backups, while another is created for differential backups (using incremental storage policy), and another for transaction log backups.
STORAGE POLICY COPIES	Data Verification	✓	
	Job Based Pruning	✓	
	Manual Retention	✓	
	Mark Job Disabled	✓	
	Selective Copy	✓	
	Comments		
SUBCLIENT POLICIES	SubClient Policy		
	Comments		
UPGRADE	Netware - Local		
	Unix - Remote (Push)		
	Unix/Linux/Macintosh - Local		
	Unix/Linux/Macintosh - Silent		
	Upgrade from CommCell Console	✓	
	Windows - Local	✓	
	Windows - Remote (Push)	✓	
	Windows - Silent	✓	
	Comments	✓	Upgrade the Agent - Non-Microsoft Clustered Environment on Windows
USER ADMINISTRATION AND SECURITY	Backup Set/Archive Set		

	Subclient		
	Comments		

Additional features are listed below:

Activity Control	Auxiliary Copy
CommCell Console	Deconfiguring Agents
GridStor	Languages
Log Files	MediaAgent
Operation window	QR Volume Creation Options
Robust Network Layer	Scheduling
SnapProtect Backup	Snapshot Engines
VaultTracker Enterprise	VaultTracker
Report Output Options	Restore/Recover/Retrieve - Other Options
Cloud Storage	Job Restart - Data Collection

Getting Started Deployment - SQL Server iDataAgent

◀ Previous Next ▶

WHERE TO INSTALL

Install the software on a computer on which SQL Server resides, and satisfies the minimum requirements specified in the System Requirements.

INSTALLATION

The software can be installed using one of the following methods:

METHOD 1: INTERACTIVE INSTALL

Use this procedure to directly install the software from the installation package or a network drive.

METHOD 2: INSTALL SOFTWARE FROM COMMCELL CONSOLE

Use this procedure to install remotely on a client computer.

METHOD 1: INTERACTIVE INSTALL

- Log on to the client computer as Administrator or as a member of the Administrator group on that computer.
- Run **Setup.exe** from the **Software Installation Package**.

If you are installing on Windows Server Core editions, navigate to Software Installation Package through command line, and then run **Setup.exe**.
- Select the required language.
Click **Next**.
- Select the option to install software on this computer.

The options that appear on this screen depend on the computer in which the software is being installed.
- Select **I accept the terms in the license agreement**.
Click **Next**.



BEFORE YOU BEGIN

Download Software Packages

Download the latest software package to perform the install.

Verify System Requirements

Make sure that the computer in which you wish to install the software satisfies the System Requirements.



- Expand **Client Modules | Backup & Recovery | Database** and then click **SQL Server iDataAgent** box.

Click **Next**.

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

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

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





- Enter the fully qualified domain name of the **CommServe Host Name**.

Click **Next**.

Do not use space and the following characters when specifying a new name for the CommServe Host Name:

`\\`~!@#$$%^&*()+=<>/?,[\]{:;'"`

- Click **Next**.



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

Click **Next**.

This option enables CommCell operations across Windows firewall by adding CommCell programs and services to Windows firewall exclusion list.

It is recommended to select this option even if Windows firewall is disabled. This will allow the CommCell programs and services to function if the Windows firewall is enabled at a later time.

11. Verify the default location for software installation.

Click **Browse** to change the default location.

Click **Next**.

- Do not install the software to a mapped network drive.
- Do not use the following characters when specifying the destination path:

/ : * ? " < > | #

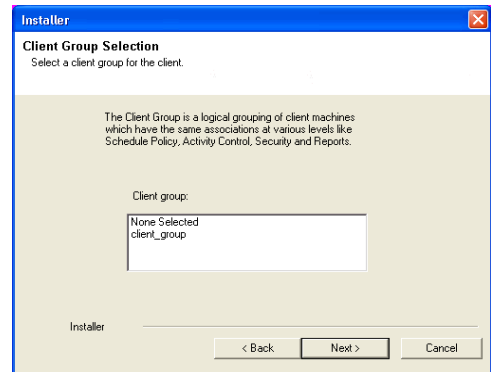



It is recommended that you use alphanumeric characters only.

12. Select a Client Group from the list.


Click **Next**.

This screen will be displayed if Client Groups are configured in the CommCell Console.


13. Click **Next**.



14. Click **Next**.




15. Select a storage policy from the **Storage Policy** list.
Click **Next**.




If you do not have Storage Policy created, this message will be displayed.
Click **OK**.

You can create the Storage Policy later in step 20.

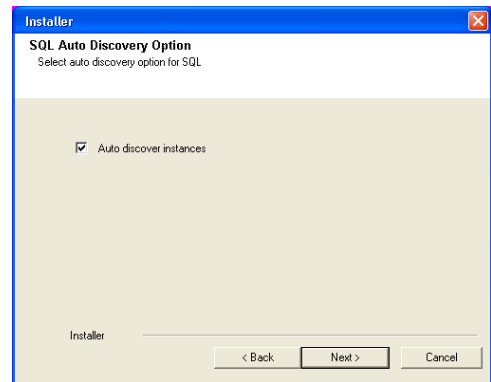


16. Click **Next**.

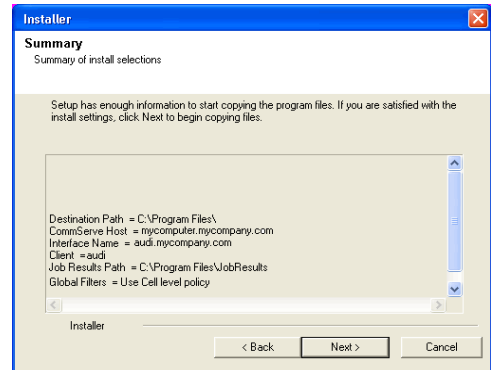
When **Auto Discover Instances** is enabled, new instances are automatically discovered every 24 hours.



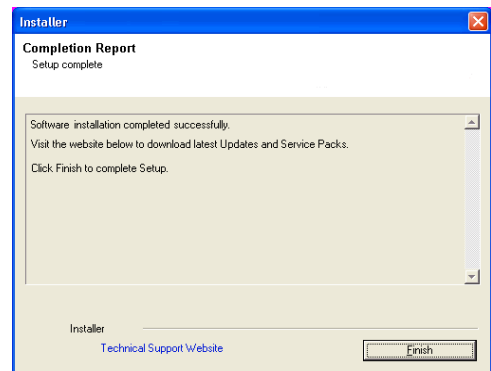
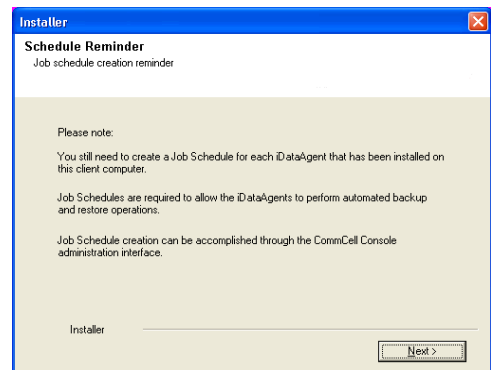
17. Click **Next**.



18. Click **Next**.



19. Click **Finish**.



If you already have a storage policy selected in step 15, Click **Next** ► button available at the bottom of the page to continue.

If you do not have Storage Policy created, continue with the following step.

- 20 To create a storage policy, you must have configured a library in the CommCell.
- If you do not already have a library configured, go to Disk Library Creation.
 - If you have a library configured, go to Storage Policy Creation.

DISK LIBRARY CREATION:

1. From the CommCell Console, click the **Backup Target** button on **EZ Operations Wizard**.
2. Click **Disc Library (For backup to disc)** and click **Next**.
3. Click **Use Local Disk**.


Type the name of the folder in which the disc library must be located in the **Enter backup destination folder** box or click the **Browse** button to select the folder.

Click **Next**.

If you click the **Use Network Share** option you will be prompted for the credentials (user name and password) to access the share.

4. Click **Next**.
5. Click **Finish**.

This will create a library and Storage Policy. Click the **Next** button available at the bottom of the page to continue.



STORAGE POLICY CREATION

1. From the CommCell Browser, navigate to **Policies**.
2. Right-click the **Storage Policies** and then click **New Storage Policy**.
3. Follow the prompts displayed in the Storage Policy Wizard. The required options are mentioned below:

- o Select the Storage Policy type as **Data Protection and Archiving** and click **Next**.
- o Enter the name in the **Storage Policy Name** box and click **Next**.
- o From the **Library** list, click the name of a disk library to which the primary copy should be associated and then click **Next**.


Ensure that you select a library attached to a MediaAgent operating in the current release.

- o From the **MediaAgent** list, click the name of a MediaAgent that will be used to create the primary copy and then click **Next**.
- o For the device streams and the retention criteria information, click **Next** to accept default values.
- o Select **Yes** to enable deduplication for the primary copy.
- o From the **MediaAgent** list, click the name of the MediaAgent that will be used to store the Deduplication store.

Type the name of the folder in which the deduplication database must be located in the Deduplication Store Location or click the Browse button to select the folder and then click **Next**.

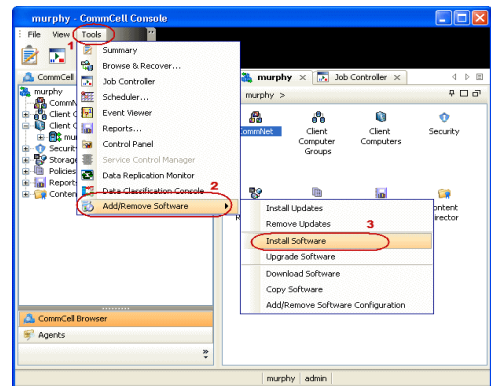
- o Review the details and click **Finish** to create the Storage Policy.

This will create a storage policy. Click the **Next** button available at the bottom of the page to continue.



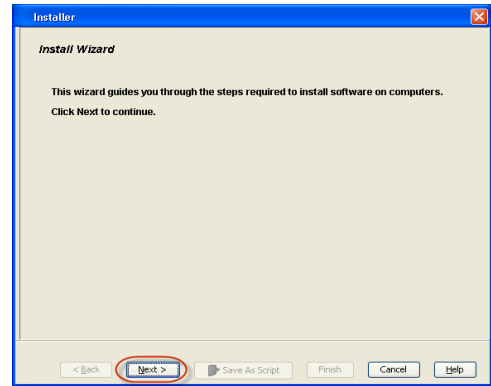
METHOD 2: INSTALL SOFTWARE FROM COMMCELL CONSOLE

1. From the CommCell Browser, select **Tools | Add/Remove Software | Install Software**.




2. Click **Next**.


3. Select **Windows**.
Click **Next**.




4. Select **Manually Select Computers**.
Click **Next**.



5. Enter the fully qualified domain name of the computer on which SQL Server resides.
For example: apple.domain.company.com
Click **Next**.



6. Click **Next**.



- Specify **User Name** and **Password** that must be used to access the client computer.
Click **Next**.

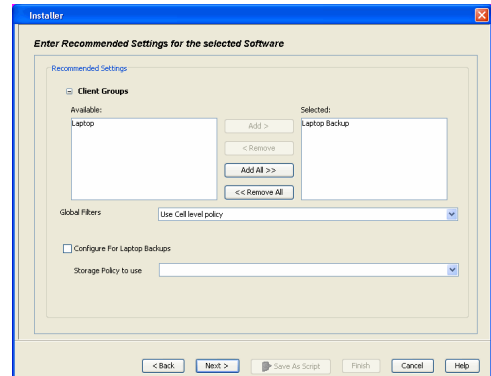



The user must be an Administrator or a member of the Administrator group on that computer.

- Select **SQL Server iDataAgent**.
Click **Next**.

- Select **Client Group** from **Available** and click **Add**.
 - From **Storage Policy to use** list, click storage policy.
 - Click **Next**.

- Click **Next**.

When **Auto Discover Instances** is enabled, new instances are automatically discovered every 24 hours.






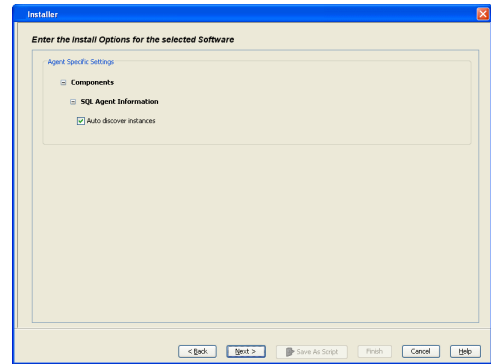
11. Click **Reboot (if required)** and then click **Next**.


When **Reboot (if required)** is selected, the install program will automatically reboot the client computer if a reboot is required during installation.

12. Click **Immediate**.
Click **Next**.

13. Click **Finish**.

14. You can track the progress of the job from the **Job Controller** or **Event Viewer** window.





ADDITIONAL INSTALLATION METHODS

Custom Package

Create a compact software package for quick deployment to multiple clients.

Decoupled Install

Install the software first and later register the client in the CommCell.

Remote Install

Deploy the software from CommCell Console on multiple clients.

Installing Restore only Agents

Setup a client in the CommCell for restore purposes.

Silent Install

Deploy the software silently on multiple clients.

Getting Started Deployment on a Cluster - SQL Server iDataAgent

◀ Previous Next ▶

SKIP THIS PAGE IF YOU ARE NOT INSTALLING THIS AGENT ON A CLUSTER.

Click **Next** ▶ to continue with the deployment.

WHERE TO INSTALL

Install the software from the active node in the cluster group. Make sure the cluster group satisfies the minimum requirements specified in the System Requirements.

1. Log on to the active node as the Domain User with administrative privileges to all nodes on the cluster.
2. Run **Setup.exe** from the **Software Installation Package**.
3. Select the required language.
Click **Next**.

4. Select the option to install software on this computer.
The options that appear on this screen depend on the computer in which the software is being installed.

5. Select **I accept the terms in the license agreement**.
Click **Next**.

6. Select **Configure a Virtual Server**.
Click **Next**.

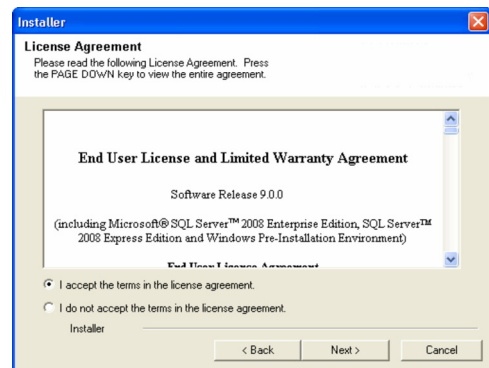


BEFORE YOU BEGIN

Download Software Packages


Download the latest software package to perform the install.

Verify System Requirements


Make sure that the computer in which you wish to install the software satisfies the System Requirements.



7. Select the disk group in which the cluster group resides.
Click **Next**.



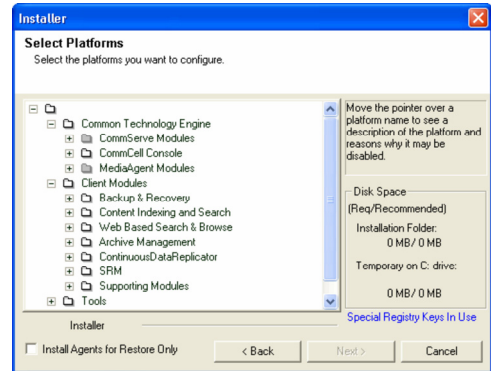
8. Expand **Client Modules | Backup & Recovery | Database** and select **SQL Server iDataAgent**
Click **Next**.



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

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

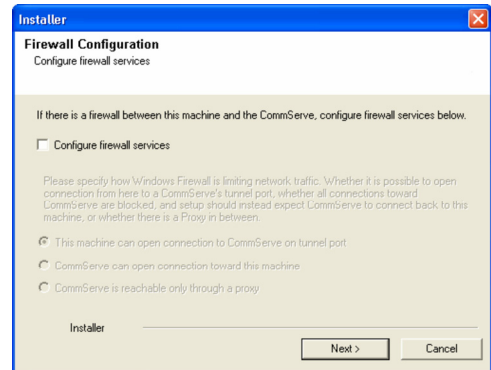
If firewall configuration is not required, click **Next**.




10. Enter the fully qualified domain name of the **CommServe Host Name**.
Click **Next**.

Do not use space and the following characters when specifying a new name for the CommServe Host Name:

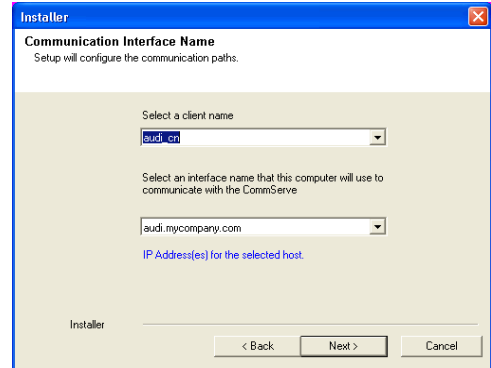
`|\`~!@#%&*()+=<>/?,[\{\}::;"`



11. Specify the name of the **Virtual Machine**.
Click **Next**.




12. Select the name of the **Physical Machine** from drop-down list.
Click **Next**.



13. Select **Add programs to the Windows Firewall Exclusion List**, to add CommCell programs and services to the Windows Firewall Exclusion List.
Click **Next**.

This option enables CommCell operations across Windows firewall by adding CommCell programs and services to Windows firewall exclusion list.

It is recommended to select this option even if Windows firewall is disabled. This will allow the CommCell programs and services to function if the Windows firewall is enabled at a later time.




14. Type or **Browse** to specify the software installation path which must be located on local disk of your physical machine and then click **Next**.

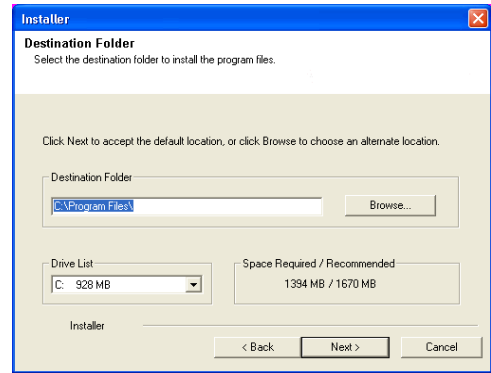
- Do not install the software to a mapped network drive.
- Do not use the following characters when specifying the destination path:

/ : * ? " < > | #

It is recommended that you use alphanumeric characters only.

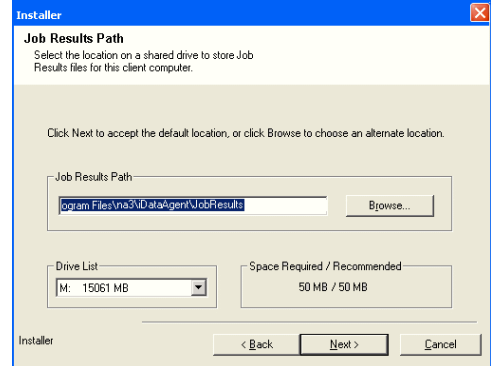


15. Verify the default location for Job Results Path.
Click **Browse** to change the default location.
Click **Next**.

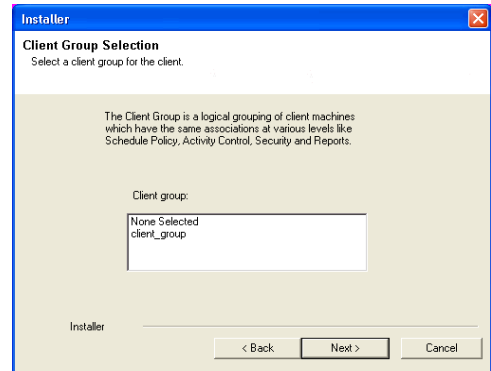


16. Select a Client Group from the list.
Click **Next**.


This screen will be displayed if Client Groups are configured in the CommCell Console.



17. Click **Next**.



18. Click **Next**.



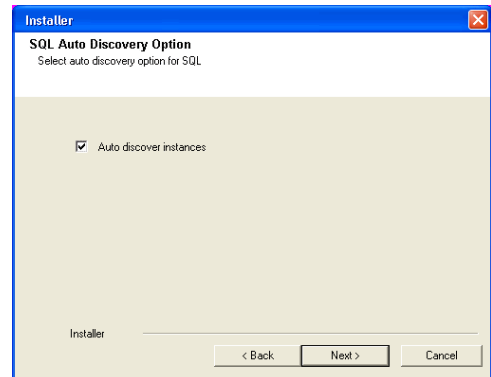
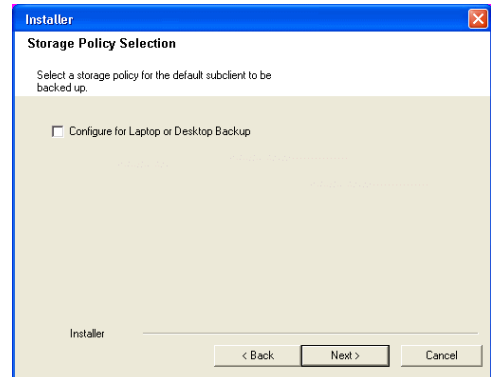


- 19. Select a **Storage Policy** from the drop-down list.
Click **Next**.

If you do not have Storage Policy created, this message will be displayed.
Click **OK**.

You can create the Storage Policy later in step 29.

- 20. Click **Next**.
For cluster, you need to discover instances manually. See Manually Discovering New Instances to discover the new instances.

- 21. Click **Next**.



22. Click **Yes**.

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

Once you complete your selections, click **Next**.

- 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.

24. Specify **User Name** and **Password** for the **Domain Administrator account Information** to perform the remote install on the cluster nodes you selected in the previous step.

Click **Next**.

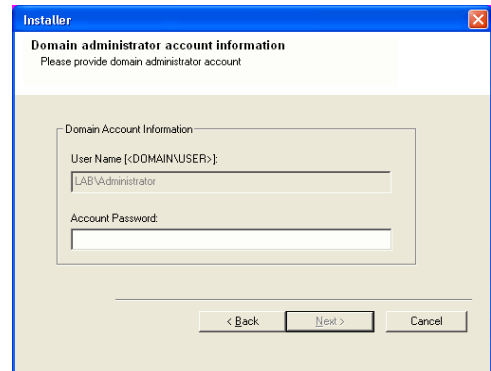



25. The progress of the remote install for the cluster nodes is displayed; the install can be interrupted if necessary.

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

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

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

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



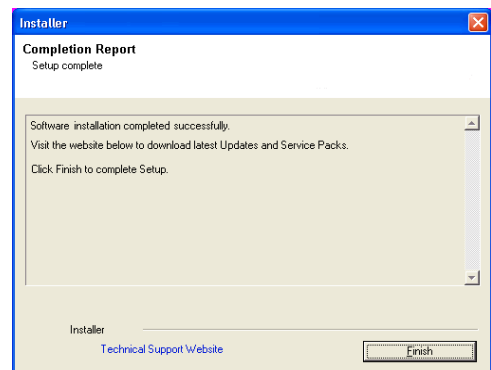



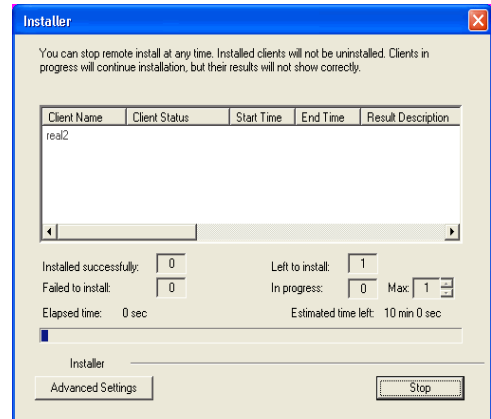
26. Read the summary for remote installation to verify that all selected nodes were installed successfully.

Click **Next**.

- 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.

27. Click **Next**.

28. Click **Finish**.





If you already have a storage policy selected in step 19, Click **Next** ► button available at the bottom of the page to continue.

If you do not have Storage Policy created, continue with the following step.

29 To create a storage policy, you must have configured a library in the CommCell.

- If you do not already have a library configured, go to Disk Library Creation.
- If you have a library configured, go to Storage Policy Creation.

DISK LIBRARY CREATION:

1. From the CommCell Console, click the **Backup Target** button on **EZ Operations Wizard**.
2. Click **Disc Library (For backup to disc)** and click **Next**.
3. Click **Use Local Disk**.


Type the name of the folder in which the disc library must be located in the **Enter backup destination folder** box or click the **Browse** button to select the folder.

Click **Next**.

If you click the **Use Network Share** option you will be prompted for the credentials (user name and password) to access the share.

4. Click **Next**.
5. Click **Finish**.

This will create a library and Storage Policy. Click the **Next** ► button available at the bottom of the page to continue.



STORAGE POLICY CREATION

1. From the CommCell Browser, navigate to **Policies**.
2. Right-click the **Storage Policies** and then click **New Storage Policy**.
3. Follow the prompts displayed in the Storage Policy Wizard. The required options are mentioned below:
 - Select the Storage Policy type as **Data Protection and Archiving** and click **Next**.
 - Enter the name in the **Storage Policy Name** box and click **Next**.
 - From the **Library** list, click the name of a disk library to which the primary copy should be associated and then click **Next**.


Ensure that you select a library attached to a MediaAgent operating in the current release.

- From the **MediaAgent** list, click the name of a MediaAgent that will be used to create the primary copy and then click **Next**.
- For the device streams and the retention criteria information, click **Next** to accept default values.
- Select **Yes** to enable deduplication for the primary copy.
- From the **MediaAgent** list, click the name of the MediaAgent that will be used to store the Deduplication store.

Type the name of the folder in which the deduplication database must be located in the Deduplication Store Location or click the Browse button to select the folder and then click **Next**.

- Review the details and click **Finish** to create the Storage Policy.

This will create a storage policy. Click the **Next** ► button available at the bottom of the page to continue.



Getting Started Deployment On A Non-Microsoft Cluster - SQL Server iDataAgent

◀ Previous Next ▶

SKIP THIS PAGE IF YOU ARE NOT INSTALLING THIS AGENT ON A NON-MICROSOFT CLUSTER.

Click **Next** ▶ to begin Configuration.

INSTALLATION

The software can be installed in one for the following Non-Microsoft Clustered environment:

INSTALL SQL SERVER IDATAAGENT IN VERITAS CLUSTER ENVIRONMENT

Use this procedure to install the software from the installation package or a network drive on a VERITAS Cluster environment.

INSTALL SQL SERVER IDATAAGENT IN HP SCALABLE NAS/POLYSERVE CLUSTER ENVIRONMENT

Use this procedure to install the software from the installation package or a network drive on a NAS/Polyserve Clustered environment.

BEFORE YOU BEGIN

Download Software Packages

Download the latest software package to perform the install.

Verify System Requirements

Make sure that the computer in which you wish to install the software satisfies the System Requirements.

INSTALL SQL SERVER IDATAAGENT IN VERITAS CLUSTER ENVIRONMENT

WHERE TO INSTALL


Install the software in a VERITAS Clustered Environment, by locating the active node that is hosting the cluster resources and perform the installation from this node to the virtual node.

Follow the steps given below to install the SQL File System iDataAgent in a VERITAS Clustered environment

1. Create `bIgnoreClusterVMCheck` registry key to detect that the agent is being installed on a virtual node in a non-Microsoft cluster.
2. Run **Setup.exe** from the **Software Installation Package**.
3. Select the required language.
Click **Next**.

To create the registry key, see the following steps:

1. Start the Registry Editor (Regedit.exe or Redegt32.exe)
2. Locate the directory under which you want to create a key, e.g., `HKEY_LOCAL_MACHINE\SOFTWARE\`.
3. Right click the **SOFTWARE** and click **New** -> **Key**.
4. Name the key as **GalaxyInstallerFlags**.
5. Right click the **GalaxyInstallerFlags** and select **New** -> **DWORD** value, name it as `bIgnoreClusterVMCheck` and by double clicking the `bIgnoreClusterVMCheck` key modify the **Value** data to **1**.



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

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


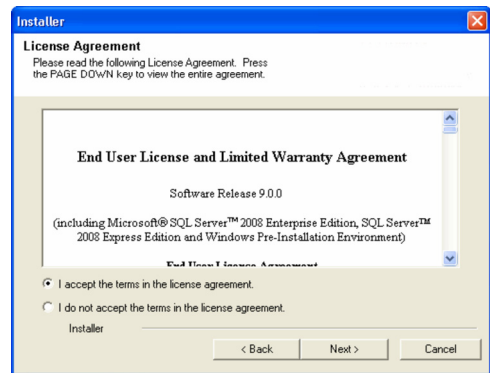

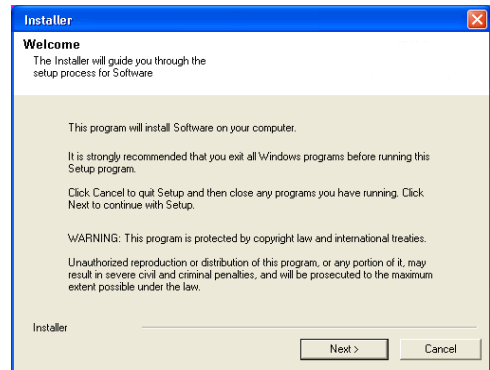

5. Click **Next**.

6. Click **OK**.

7. Select **I accept the terms in the license agreement**.
Click **Next**.

8. Select **Configure a Virtual Server**.
Click **Next**.

- 9.
- In the **disk groups** box, specify the host name of the Virtual Host created for backup, e.g., VirtualVeritasServe.
 - In the **Communication Interface Name** box, specify the fully qualified domain name of the Virtual Host, e.g., VirtualVeritasServe.acme.com.



- Click **Next**.

10. Specify the hosts node in the **Preferred Nodes** and click **Add** to add in the Nodes on which **Virtual Group is configured to run**.

Specify the name of active node in **Current Owner Node** box.

Click **Next**.

11. Click **Yes** if it is the Active node.

- If you click **No** the program will perform the passive install. Follow the steps described in Manually Installing the Software on a Passive Node.
- This screen may look different from the example shown.

12. Expand **Client Modules | Backup & Recovery | Database** and then click **SQL Server iDataAgent** box.

Click **Next**.

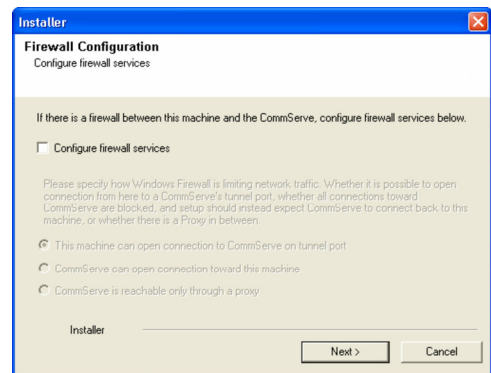
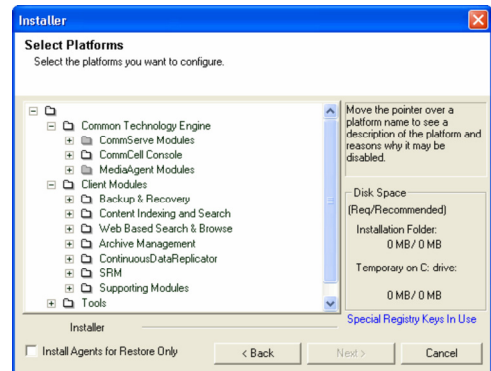


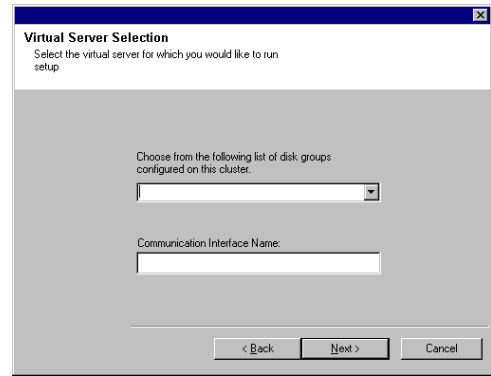
13. If this computer and the CommServe is separated by a firewall, select the **Configure firewall services** option and then click **Next**.

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

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

14. Enter the fully qualified domain name of the **CommServe Host Name**.


Click **Next**.



Do not use space and the following characters when specifying a new name for the CommServe Host Name:

\|`~!@#\$%^&*()+=<>/?,[\]{ };"

15. Click **Next**.




16. Select **Add programs to the Windows Firewall Exclusion List**, to add CommCell programs and services to the Windows Firewall Exclusion List.

Click **Next**.

This option enables CommCell operations across Windows firewall by adding CommCell programs and services to Windows firewall exclusion list.

It is recommended to select this option even if Windows firewall is disabled. This will allow the CommCell programs and services to function if the Windows firewall is enabled at a later time.



17. Verify the default location for software installation.


Click **Browse** to change the default location.

Click **Next**.

- Do not install the software to a mapped network drive.
- Do not use the following characters when specifying the destination path:

/ : * ? " < > | #

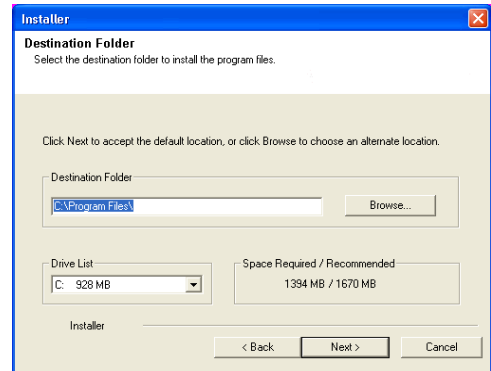
It is recommended that you use alphanumeric characters only.



18. Specify the **Job Results Path**. Make sure that the Job Results folder you specify resides on a shared disk.

Click **Browse** to change the default location.

Click **Next**.



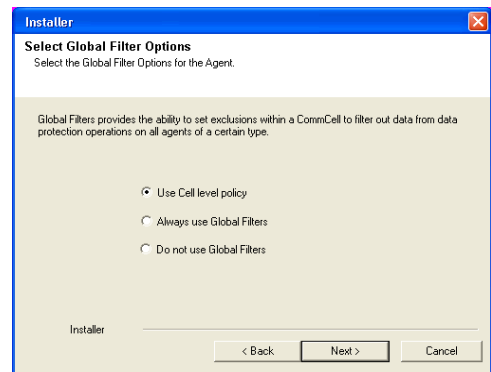
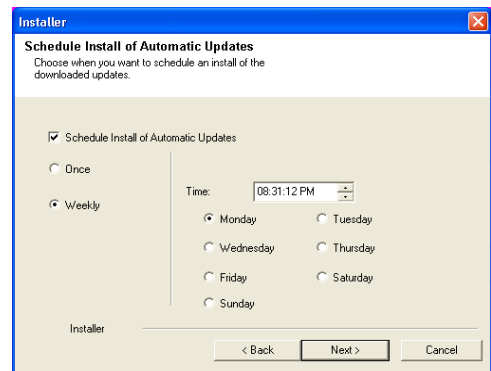
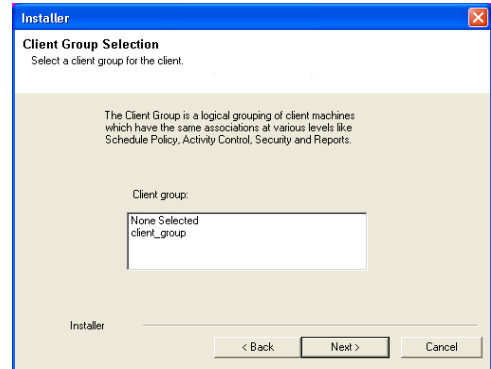
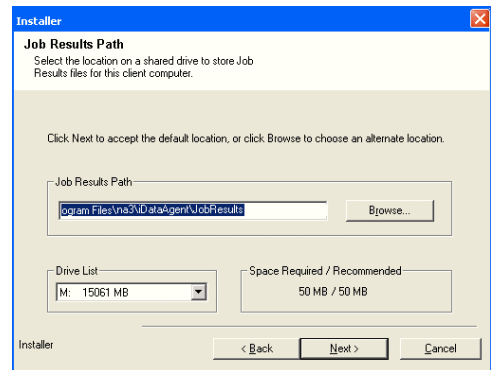
19. Select a Client Group from the list.
Click **Next**.

This screen will be displayed if Client Groups are configured in the CommCell Console.

20. Click **Next**.

21. Click **Next**.

22. Select a **Storage Policy** from the drop-down list.
Click **Next**.



23. Click **Next**.

This process will not automatically discover the Virtual nodes. It is recommended to configure the Virtual node post installation from CommCell Console. For details, see [Virtual Nodes Are Not Being Automatically Discovered in a Veritas Cluster Environment](#)

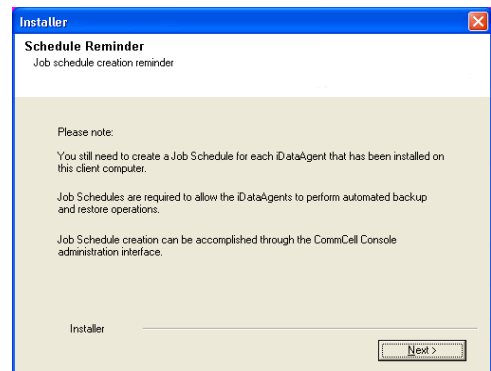

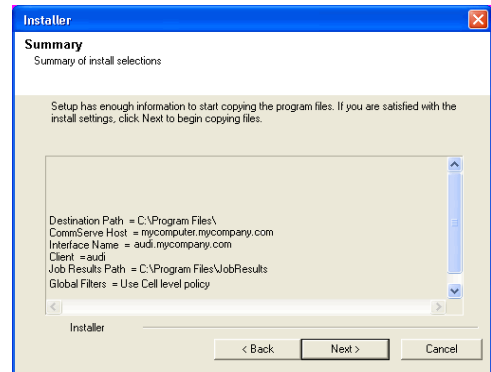


24. Click **Next**.


25. Click **OK** and manually configure the cluster resources.

See [Configuring the VERITAS Cluster Resources](#) for more information on configuring the cluster resources.

26. Click **Next**.

27. Click **Finish**.





CONFIGURING VERITAS CLUSTER RESOURCES

In case of VERITAS Cluster you must manually install the software on all the available Nodes. For installs on the VERITAS Cluster, complete the following procedures for the agent. These procedures can be completed using the VERITAS Cluster Administrator.

CREATE THE CLUSTER RESOURCES FOR SERVICES

28. Once the physical nodes and the virtual machine are installed, services should be added as generic services.

Also, you may want to designate the services as "critical" if you want to cause the virtual machine to fail over if the associated service stops.

- In Windows Services, locate the services for each installed cluster (e.g., Bull Calypso Cluster Plugin (penodevn1_netwo) (instance001)) right-click and select **Properties**, and locate the service name in the **General** tab (e.g., GxClusPlugin (penodevn1_netwo) (Instance001))

The service name is case- and space-sensitive, so record it exactly as it appears.

29. Create Service Dependencies

- In Cluster Administrator, navigate to the appropriate Cluster group. Create a generic service resource for the cluster plugin service.
- If you want an automatic failover to occur, mark the resource as **Critical** and **Enabled**.
- Repeat this procedure for each client installed on the cluster.

30. Associate Services with the Current Active Node

- In Cluster Administrator, click **Resources**.
- Click **Link** and make the Cluster plugin resource dependent upon the **MountV** resource and the **Lanman** service.
- In Cluster Administrator, navigate to the appropriate **GenericService** group.
- Right-click the cluster plugin resource.
- From the short-cut menu, click **Online** and the name of the current active node.

INSTALL SQL SERVER IDATAAGENT IN HP SCALABLE NAS/POLYSERVE CLUSTER ENVIRONMENT

WHERE TO INSTALL

Install the software in HP Scalable NAS/Polyserve Cluster environment from the active host to virtual server.

The software must be installed on each passive node available to the virtual host. During install, you will be asked to provide the name of the active host for each passive host installation. See Manually Installing the Software on a Passive Node for step-by-step instructions.

Follow the steps given below to install the Windows File System iDataAgent in a HP Scalable NAS/Polyserve Clustered environment:


1. Create `bIgnoreClusterVMCheck` registry key to detect that the agent is being installed on a virtual node in a non-Microsoft cluster.

To create the registry key, see the following steps:


1. Start the Registry Editor (Regedit.exe or Regedt32.exe)
2. Locate the directory under which you want to create a key, e.g., `HKEY_LOCAL_MACHINE\SOFTWARE\`.
3. Right click the **SOFTWARE** and click **New -> Key**.
4. Name the key as **GalaxyInstallerFlags**.
5. Right click the **GalaxyInstallerFlags** and select **New -**

> **DWORD** value, name it as `bIgnoreClusterVMCheck` and by double clicking the `bIgnoreClusterVMCheck` key modify the **Value** data to **1**.

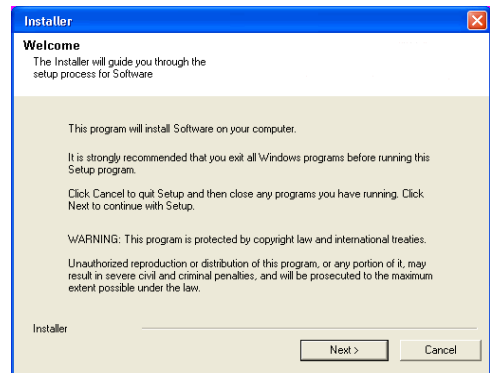
2.
 - In the PolyServe Matrix Server, create a virtual host resource name that has an assigned TCP/IP address in the network.
 - Provide a virtual host name (e.g., VirtualPolyServe.acme.com) and an Application name (e.g., VirtualPolyServe).
 - Select all network Interfaces that will be available as hosts for the backup software services and click **OK**.
3. From the Virtual Hosts tab in the PolyServe Matrix Server dialog box, locate the primary node for the virtual server you are installing, and perform the installation from the active (primary) host to the virtual server.
4. Run **Setup.exe** from the **Software Installation Package**.
5. Select the required language.
Click **Next**.




6. Select the option to install software on this computer.
The options that appear on this screen depend on the computer in which the software is being installed.



7. Click **Next**.




8. Click **OK**.




9. Select **I accept the terms in the license agreement**.
Click **Next**.


10. Select **Configure a Virtual Server**.
Click **Next**.



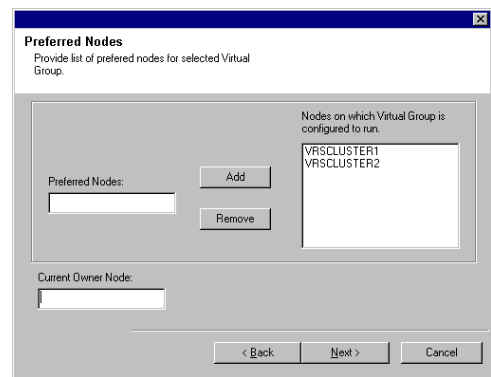
11.
 - In the **disk groups** box, specify the host name of the Virtual Host created for backup, e.g., VirtualVeritasServe.
 - In the **Communication Interface Name** box, specify the fully qualified domain name of the Virtual Host, e.g., VirtualVeritasServe.acme.com.
 - Click **Next**.




12. Specify the hosts node in the **Preferred Nodes** and click **Add** to add in the Nodes on which **Virtual Group is configured to run**.
Specify the name of active node in **Current Owner Node** box.
Click **Next**.



13. Click **Yes** if it is the Active node.
 - If you click **No** the program will perform the passive install. Follow the steps described in Manually Installing the Software on a Passive Node.
 - This screen may look different from the example shown.



14. Expand **Client Modules | Database** and then click **SQL Server iDataAgent** box.
Click **Next**.



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

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

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

16. Enter the fully qualified domain name of the **CommServe Host Name**.

Click **Next**.

Do not use space and the following characters when specifying a new name for the CommServe Host Name:

`\|`~!@#$%^&*()+=<>/?,[\{\};:;''`

17. Click **Next**.





The client name should be equivalent to the vhost name and the interface name equals the vhosts fully qualified domain name.

18. Select **Add programs to the Windows Firewall Exclusion List**, to add CommCell programs and services to the Windows Firewall Exclusion List.

Click **Next**.

This option enables CommCell operations across Windows firewall by adding CommCell programs and services to Windows firewall exclusion list.

It is recommended to select this option even if Windows firewall is disabled. This will allow the CommCell programs and services to function if the Windows firewall is enabled at a later time.



19. Verify the default location for software installation.

Click **Browse** to change the default location.

Click **Next**.

- 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.

20. Specify the **Job Results Path**. Make sure that the Job Results folder you specify resides on a shared disk.

Click **Browse** to change the default location.

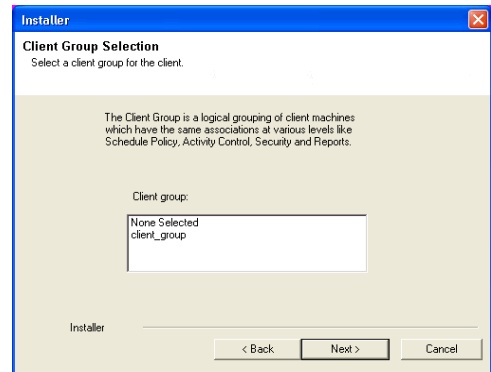


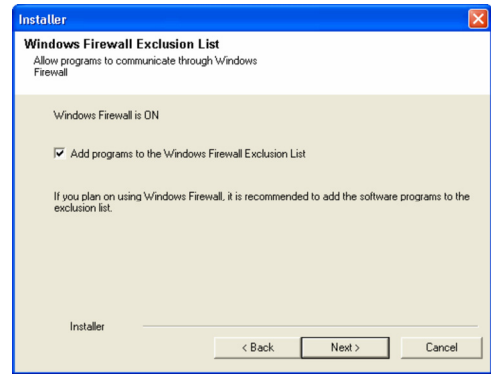
Click **Next**.

21. Select a Client Group from the list.


Click **Next**.

This screen will be displayed if Client Groups are configured in the CommCell Console.


22. Click **Next**.




23. Click **Next**.



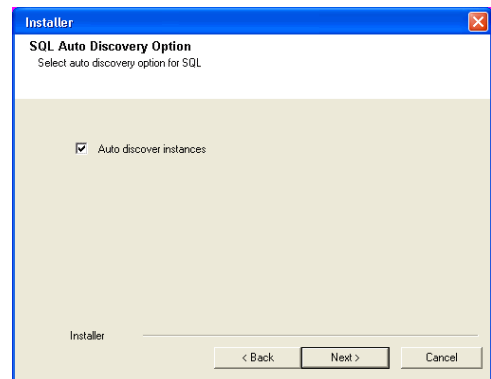
24. Select a **Storage Policy** from the drop-down list.
Click **Next**.



25. Click **Next**.
Auto Discover Instances are automatically discovered for every 24 hours.



26. Click **Next**.


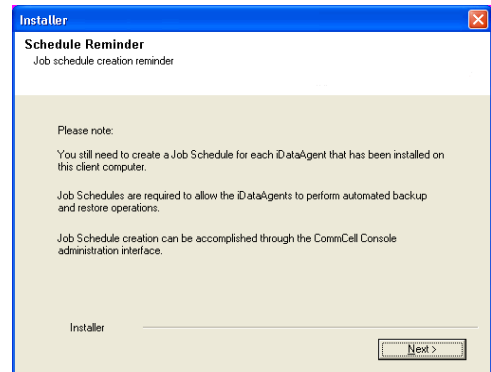




27. Click **OK** and manually configure the cluster resources.

See [Configuring the HP Scalable NAS/Polyserve Cluster Resources](#) for more information on configuring the cluster resources.

28. Click **Next**.

29. Click **Finish**.



CONFIGURING HP SCALABLE NAS/POLYSERVE CLUSTER RESOURCES

For installs on the HP Scalable NAS/PolyServe Cluster, complete the following procedures for the agent.

30. Create and Use Scripts to Administer Virtual Machine Services.

You should create and use three scripts to administer virtual machine services:

- Probe script is used to reveal the state of the specified services. This is done using the `GxSCMcmd.exe` utility.
- Start script is used to start services.
- Stop script is used to stop services.

Create a probe script to have the status of both the Bull Calypso Communications Service (GxCVD) and the Bull Calypso Event Manager Service (GxEvMgrC) checked.

In the probe script, include the `GxSCMcmd.exe` utility by name in the appropriate command. (This utility is available from the Resource Pack. See [Tools and Utilities](#) for more information.)

Create a folder (e.g., m:\GxSCMcmd) on a shared disk that the virtual host you created has access to.

Copy the GxSCMcmd.exe utility and the script files you created (probe.bat, start.bat, stop.bat) to this folder.

The utility will return the state of the specified services: '0' if the services are running or '1' if services are stopped. For an error condition, '-1' is returned.

Note that the service name includes the name of the client and the name of the instance. To avoid error in providing the vhost name and instance name in the path of scripts, copy the key name from here:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\GXClusPlugIn
(vhostname) (Instance00x)
```

These examples provide sample script names, locations and virtual host name as suggested names not requirements.

The command syntax is:

```
GxSCMcmd.exe -IsRunning <list of service name>
```

For example, your probe script (e.g., m:\GxSCMcmd\gxprobe.bat) may appear as follows:

```
@echo off
if '%MX_ACTIVE_STATE%'='INACTIVE' exit 0
m:\GxSCMcmd\GxSCMcmd.exe -IsRunning "GxClusPlugIn (virtualpolyserve)
(Instance001)"
if ERRORLEVEL 0 exit 0
exit 1
```

Be sure to copy the GxSCMcmd.exe utility to the shared disk resource. Also, be sure to include the probe script on the shared disk.

The stop and start scripts must also contain the registry entry for the virtual machine service names. Your start script (e.g., m:\GxSCMcmd\gxstart.bat) may appear as follows:

```
@echo off
C:\WINDOWS\system32\net.exe start "GxClusPlugIn (virtualpolyserve)
(Instance001)"
exit 0
```

Your stop script (e.g., m:\GxSCMcmd\gxstart.bat) may appear as follows:

```
@echo off
C:\WINDOWS\system32\net.exe stop GxClusPlugIn (virtualpolyserve)
(Instance001)
exit 0
```

31. Add the Customer Service Monitor

To add the custom service monitor to the virtual host configuration in the PolyServe Matrix Server console.

1. Ensure that the virtual host name <virtualpolyserve> is online.
2. From the Services MMC snap-in, verify or manually start the cluster plugin service, "Bull Calypso Cluster Plugin (virtualpolyserve) (Instance00x)", on the active primary node for the Vhost virtualpolyserve
3. From the PolyServe Matrix Server, click the **Virtual Hosts** tab.
4. Right click **Virtual Host** <virtualpolyserve> and then **Add Service Monitor**.
5. For the **Monitor Type** field in the Add Service Monitor dialog box, click **CUSTOM**.
6. Assign a name to the service monitor by typing the name in the **Name** field.
7. In the **User probe script** field, type the path to the probe script that calls the GxSCMcmd.exe command (on the share disk)
8. Click **Advanced**.
9. From the Advanced Service Configuration dialog box, click the **Scripts** tab.
10. In the **Script pathname: Start** field, type the path to the start script (shared disk). In the corresponding **Timeout (seconds)** field, type a timeout value greater than 60 seconds.

11. In the **Script pathname: Stop** field, type the path to the stop script (shared disk). In the corresponding **Timeout (seconds)** field, type a timeout value greater than 60 seconds.
12. Click **OK** once, then again.

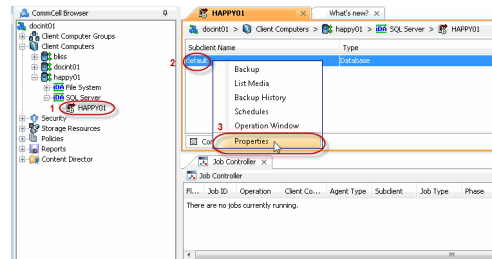
The monitor should now start, and the primary node should now be active. Failover will occur if services are disabled on the primary node, and services in such a case will move to a secondary node.



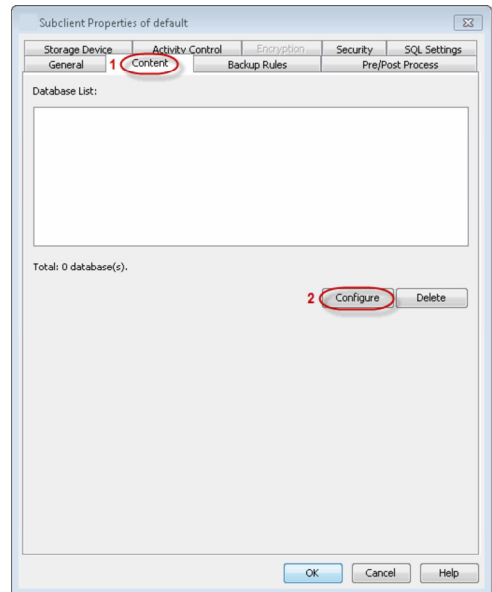
Getting Started Configuration - SQL Server iDataAgent

Once installed, configure a subclient and associate a storage policy as follows:


1.
 - From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server | <Instance>**.
 - Right-click the default subclient and then click **Properties**.



2.
 - Click the **Content** tab.
 - Click **Configure**.



3. Click **Discover**.



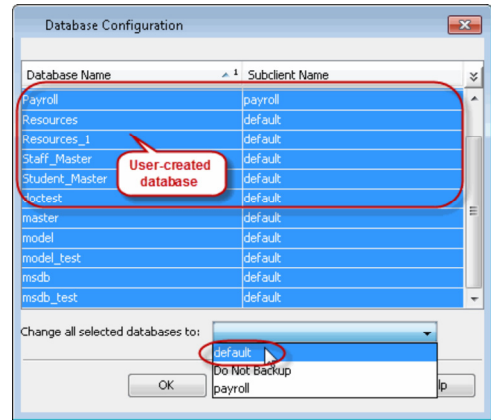
4.
 - Hold down CTRL and click each database listed in the **Database Name** list.

Ensure you have at least one user-created database selected as depicted in the image on the right.

- In the **Change all selected databases to** list, click **default**.
- Click **OK**.

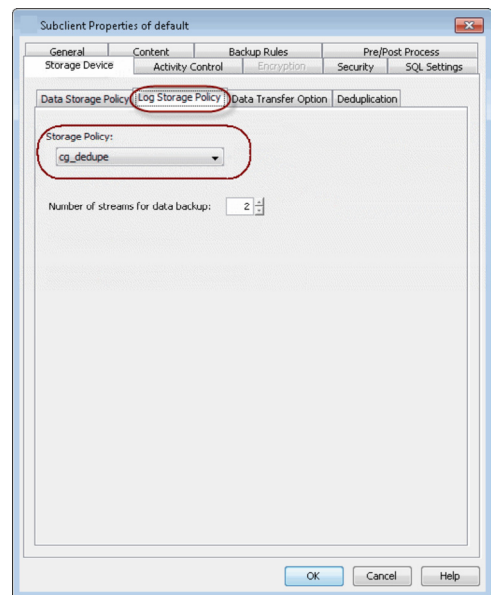
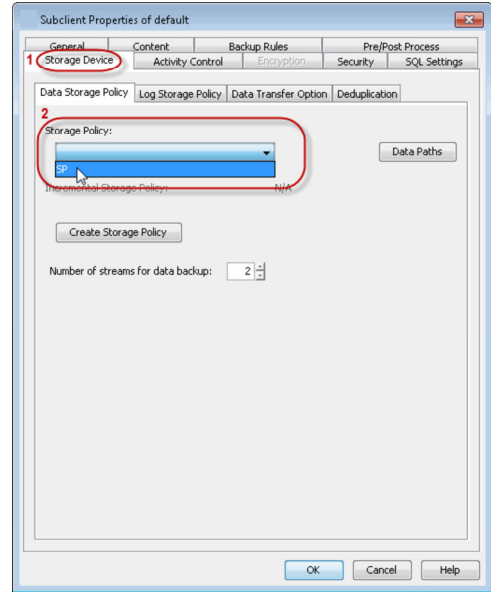
5.
 - Click the **Storage Device** tab.
 - In the **Storage Policy** box, select a Storage Policy name.

If you do not have Storage Policy created, follow the step given in step 7 below to create a storage policy.



6.
 - Click the **Log Storage Policy** tab.
 - In the **Storage Policy** box, select a Storage Policy name.
 - Click **OK**.

Click **Next >** to continue.



7. Create a Storage Policy:
 1. From the CommCell Browser, navigate to **Policies**.
 2. Right-click the **Storage Policies** and then click **New Storage Policy**.
 3. Follow the prompts displayed in the Storage Policy Wizard. The required options

are mentioned below:


- Select the Storage Policy type as **Data Protection and Archiving** and click **Next**.
- Enter the name in the **Storage Policy Name** box and click **Next**.
- From the **Library** list, click the name of a disk library to which the primary copy should be associated and then click **Next**.

Ensure that you select a library attached to a MediaAgent operating in the current release.

- From the **MediaAgent** list, click the name of a MediaAgent that will be used to create the primary copy and then click **Next**.
- For the device streams and the retention criteria information, click **Next** to accept default values.
- Select **Yes** to enable deduplication for the primary copy.
- From the **MediaAgent** list, click the name of the MediaAgent that will be used to store the Deduplication store.

Type the name of the folder in which the deduplication database must be located in the Deduplication Store Location or click the Browse button to select the folder and then click **Next**.


- Review the details and click **Finish** to create the Storage Policy.




Getting Started Backup - SQL Server iDataAgent

WHAT GETS BACKED UP	WHAT DOES NOT GET BACKED UP
<p>SQL Server system databases, which includes master, msdb, and model</p> <p>SQL Server user databases</p> <p>SQL Server Filestream Databases using regular and IntelliSnap Backups</p>	<p>tempdb - as SQL Server re-creates tempdb every time the server is started, it is not included in the backup</p> <p>SQL analysis server database</p> <p>Databases on removable media</p> <p>Databases in single-user mode, suspect, loading, standby, and offline.</p> <p>SQL Server Filestream Databases using VSS</p> <p>The following cannot be backed up using the SQL Server iDataAgent:</p> <ul style="list-style-type: none"> • Full text indexes • SQL Server application files residing on a local hard drive • File system data that resides on a SQL Server computer <p style="text-align: center;">Use the Windows File System iDataAgent to back up the above mentioned components.</p>


- From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server | <Instance>**.
 - Right-click the default subclient and click **Backup**.




- Click **Full** as backup type and then click **Immediate**.
 - Click **OK**.




- You can track the progress of the job from the **Job Controller** window of the CommCell console.



- Once the job is complete, view the job details from the **Backup History**. Right-click the **Subclient** and select **Backup History**.




- Click **OK**.



6. Right-click the job to:

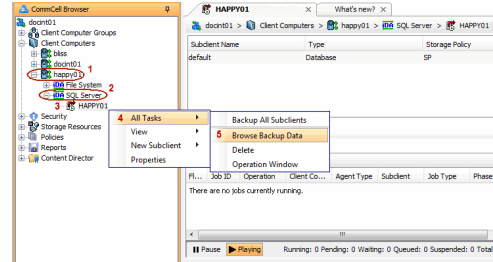
- Browse the databases that were backed up.
- View items that failed, if any, during the job.
- Resubmit the job.
- View job details.
- View media associated with the job.
- View events associated with the job.
- View backup items (you can view the database files that were backed up e.g., .mdf, .ldf).
- View or send the log file that is associated with the job.



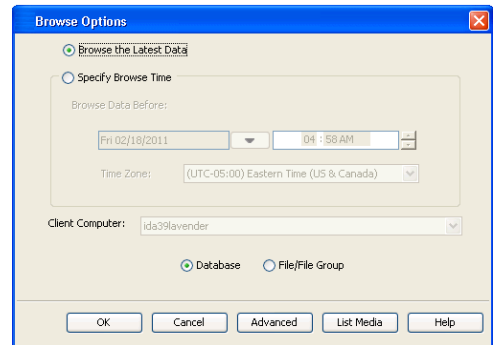
Getting Started Restore - SQL Server iDataAgent

As restoring your backup data is very crucial, it is recommended that you perform a restore operation immediately after your first full backup to understand the process. The following sections explain the steps for restoring a user-created database to a different location on the same destination server.


- From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
 - Right-click the instance and then click **All Tasks | Browse Backup Data**.



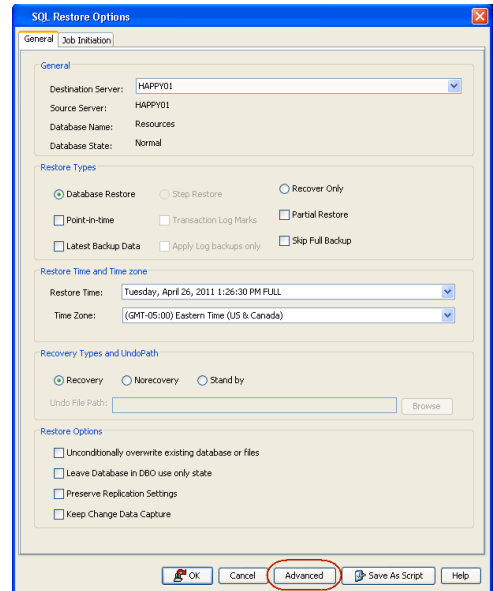
- Click **OK**.



- In the right pane of the Browse window, select the database you want to restore.
 - Click **Recover All Selected**.




- Click **Advanced**.




- Click the **Options** tab.
 - Rename the database name under the **Database** column.
 - Change the path of the database and log files under the **Physical Path** column.
 - Click **OK**.


6. Click **OK**.



7. You can monitor the progress of the restore job in the **Job Controller**.



8. Once the restore job has completed, right-click the agent and click **View | Restore History**.

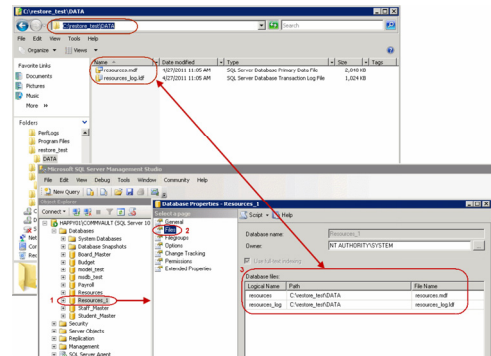

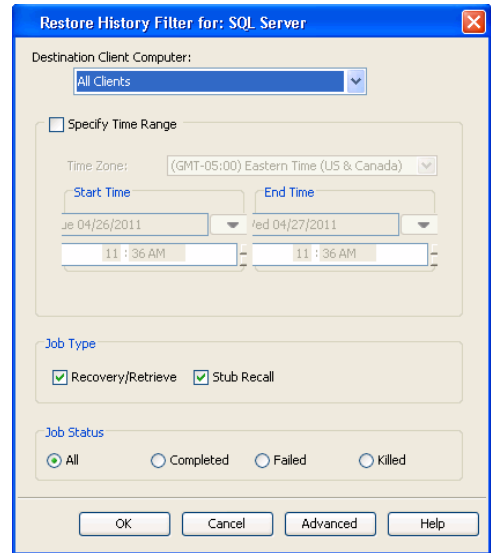


9. Click **OK**.

10. You can view the following details about the job by right-clicking the job:

- View Restore Items
 - You can view them as **Successful, Failed, Skipped** or **All**.
- View Job Details
- View Events of the restore job.
- View Log files of the restore job
- View Job Path
- Send Logs
- Resubmit Job

11. Once the database is restored, verify that the restored database and log files are available in the restore destination provided during step 5.



CONGRATULATIONS - YOU HAVE SUCCESSFULLY COMPLETED YOUR FIRST BACKUP AND RESTORE.

If you want to further explore this Agent's features read the **Advanced** sections of this documentation.



Advanced Configuration - SQL Server iDataAgent

TABLE OF CONTENTS

<p>Understanding the CommCell Console</p> <p>Creating User-Defined Subclients For Databases For Files and Filegroups</p> <p>Managing Instances Manually Discovering New Instances Enable / Disable Automatic Discovery Setting the Discovery Frequency</p> <p>Enabling/Disabling Automatic Database Discovery Manually Discovering Databases</p> <p>Automatically Discovering Databases in Offline States For All Clients For Individual Clients</p> <p>Excluding Databases from Backups</p> <p>Managing Databases Deleted from SQL Server In the Default Subclient In User-Defined and Default Subclients</p> <p>Specifying Databases for OnDemand Backup Selecting Backup types for On Demand Backups and Executing the Backups</p> <p>Specifying File and Filegroups for OnDemand Backup</p> <p>Setting Up Backup Conversion Rules For Default and Database Subclients For File/Filegroup Subclients</p> <p>Enhancing Performance During Backups</p> <p>Configuring the Number of Log Backups to Run Before a Full Backup</p> <p>Configuring Log Backups to Run Without Full Backups</p> <p>Performing Backups using VSS</p> <p>Configuring Data Streams</p> <p>Configuring User Accounts for Backups At the CommCell Level At the Client Computer Group Level At the Agent Level At the Instance Level</p> <p>Modifying an Agent, Instance, or Subclient</p> <p>Deleting an Agent, Instance, or Subclient Deleting an Agent Deleting an Instance Deleting a Subclient</p>	<p>Command Line Operations Log on to the CommServe Configure Instances Configure Subclients Perform Command Line Operations from Microsoft SQL Management Studio</p>
---	---

UNDERSTANDING THE COMMCELL CONSOLE

The Microsoft SQL Server iDataAgent uses the following logical entities to manage backup and restore operations from the CommCell Console.

AGENT


Facilitates SQL instance discovery.

INSTANCE

Defines the SQL Server instance to be backed up.

SUBCLIENT

Defines the SQL databases to be backed up.



CREATING USER-DEFINED SUBCLIENTS

By default, all databases within each SQL Server instance are automatically assigned to the default subclient. This subclient backs up the entire instance.

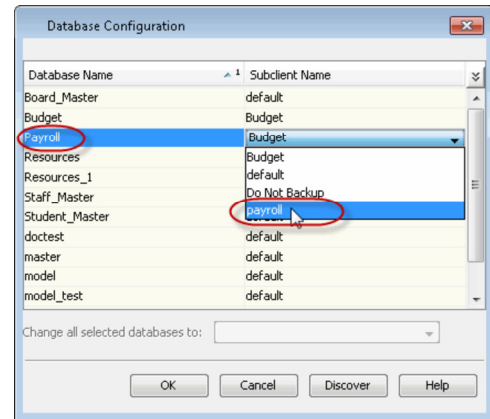
If you want to divide your backups into smaller groups, you can do so by creating user-defined subclients as described in the following sections.

FOR DATABASES

If you want to back up groups of specific databases, you can do so by creating a user-defined subclient containing any number of databases that exist within the instance. This is useful if you want to back up a subset of databases at certain times or with a particular frequency.

When you create a user-defined subclient, the contents of the user-defined subclient will be excluded from the Default Subclient.

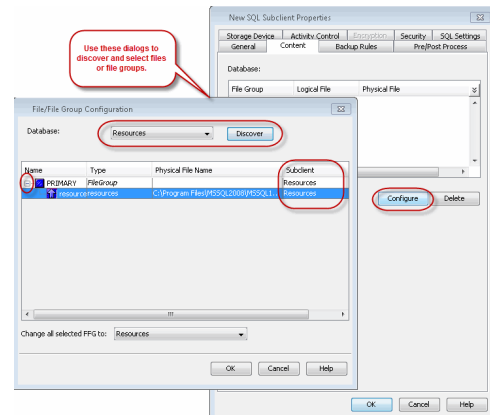
1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server | <Instance>**.
2. Right-click the **Instance**, point to **New Subclient**, and then click **Database**.
3. In the **Subclient Name** box, type a name.
4. Click the **Storage Device** tab.
5. From the **Data Storage Policy** sub-tab, click a storage policy name from the **Storage Policy** list.
6. From the **Log Storage Policy** sub-tab, click a storage policy name from the **Storage Policy** list.
7. Click the **Content** tab and then click **Configure**.
8. Click **Discover**.
9. From the **Subclient Name list** in the **Database Configuration** window, select the name of this subclient for each database you want to include.
10. Click **OK** to save the content.
11. Click **OK**.



FOR FILES AND FILEGROUPS

In many cases, large databases may contain portions of data that require more frequent backups than others. For example, tables consisting of records entered on a daily basis may require nightly backups, whereas tables consisting of records entered on a quarterly basis may require only monthly backups. You can group such elements together by creating a user-defined subclient for files or filegroups.

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server | <Instance>**.
2. Right-click the **Instance**, point to **New Subclient**, and then click **Files and Filegroups**.
3. In the **Subclient Name** box, type a name.
4. Click the **Storage Device** tab.
5. From the **Data Storage Policy** sub-tab, click a storage policy name in the **Storage Policy** list.
6. From the **Log Storage Policy** sub-tab, click a storage policy name in the **Storage Policy** list.
7. Click the **Content** tab and then click **Configure**.
8. From the **File/FileGroup Configuration** window, select the database containing the files or filegroups you want to back up from the **Database** list.
9. Click **Discover**.
10. Expand the nodes in the **Name** list.
11. In the **Subclient Name** list, select the name of this subclient for each file or filegroup you want to include.
12. Click **OK** to save the content.
13. Click **OK**.



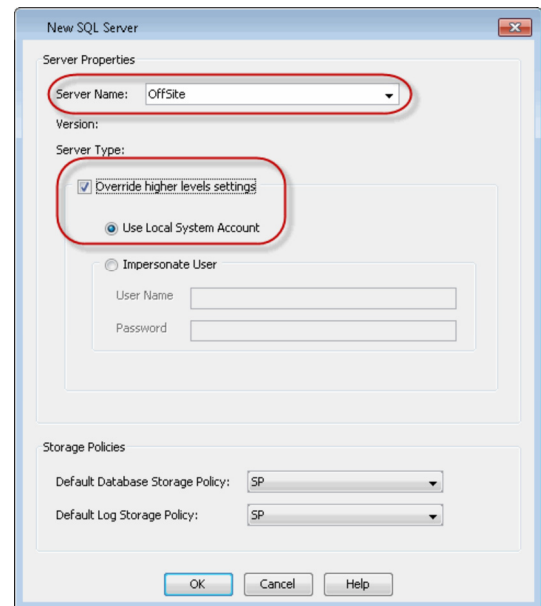
It is recommended that filegroups be configured rather than individual files. Filegroups require less overall maintenance and reduce the need to manually add or remove individual files to the subclient.

MANAGING INSTANCES

MANUALLY DISCOVERING NEW INSTANCES

By default, new instances added to the SQL Server are automatically discovered if the option to do so was enabled during the SQL Server iDataAgent installation. If this option was not enabled during installation, you can discover new instances at any time as follows:

1. Ensure you have a user account with sufficient privileges to create a new instance. Refer to the Configuring User Accounts for Backups section on this page for information on required account privileges.
2. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
3. Right-click **SQL Server** and click **New SQL Server**.
4. In the **Server Name** list, select the name of the SQL Server instance you want to assign to this instance.
5. In the **Server Type** area, check the **Override higher levels settings** check box.
6. Click **OK**.



ENABLE / DISABLE AUTOMATIC DISCOVERY

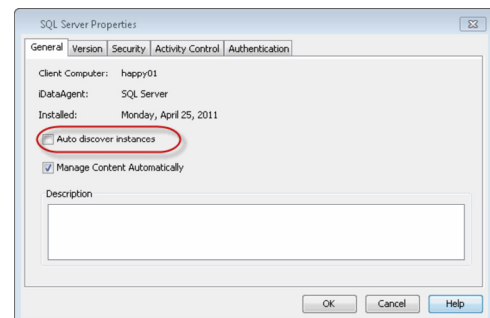
If you enable automatic discovery new SQL Server instances will be discovered as follows:

- Every 24 hours.
- Whenever the Communications Service (GxCVD) is restarted (such as after a computer reboot).

This capability ensures all instances are accounted for on a daily basis for backups.

If you want to enable or disable automatic instance discovery, you can do so as follows:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
2. Right-click **SQL Server** and click **Properties**.
3. Check or clear the **Auto discover instances** check box.
4. Click **OK**.



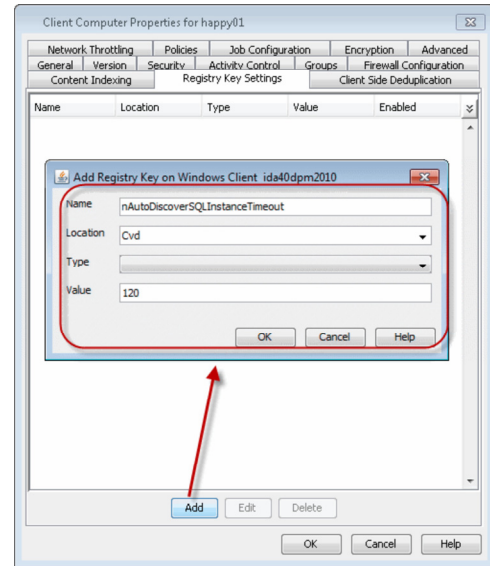
SETTING THE DISCOVERY FREQUENCY

If you want instances to be automatically discovered on a more or less frequent basis, you can do so as follows:

1. From the CommCell Browser, 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 `nAutoDiscoverSQLInstanceTimeout`.
6. In the **Location** list, select `cvd` from the dropdown list.
7. In the **Type** list, select **REG_DWORD**.
8. In the **Value** field, type the number of minutes to discover instances.

For example, to discover instances every two hours, type `120`.

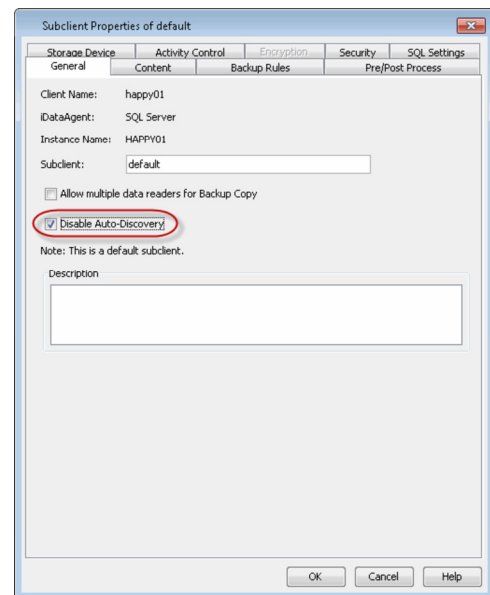
9. Click **OK**.



ENABLING/DISABLING AUTOMATIC DATABASE DISCOVERY

By default, new databases created on the SQL Server are automatically discovered and assigned to the default subclient. You can disable this functionality as follows:

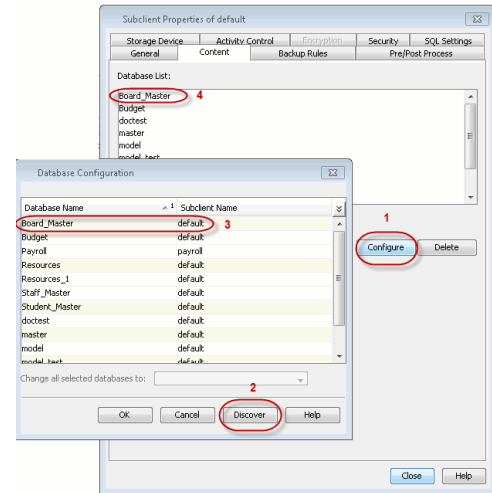
1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server | <Instance>**.
2. Right-click the default subclient and click **Properties**.
3. Enable or clear the **Disable Automatic Discovery** check box.
4. Click **OK**.



MANUALLY DISCOVERING DATABASES

If automatic discovery of databases is disabled, you can manually add databases to a subclient as follows:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server | <Instance>**.
2. Right-click the default subclient and click **Properties**.
3. Click the **Content** tab.
4. Click **Configure**.
5. Click **Discover**.
6. In the **Subclient Name** list, select the subclient to which the database you want to add should be assigned.
7. Click **OK** to save your settings.
8. Click **OK**.



AUTOMATICALLY DISCOVERING DATABASES IN OFFLINE STATES

By default, offline databases in the following states are not automatically discovered:

- Standby
- Restoring
- Suspect (will also include the Standby state)
- Shutdown (will also include the Suspect and Standby states)
- Offline

You can configure automatic discovery of offline databases for one or all clients as described below.

FOR ALL CLIENTS

1. Log on to the CommServe computer.
2. From the command prompt, navigate to `<software_installation_path>\base`.
3. Run the following command:

```
operation execscript -sn SetKeyIntoGlobalParamTbl.sql -si DatabaseStateString -si y -si <database_state>
```

where `<database_state>` is the database state(s) to be discovered.

If entering more than one state, separate each with a semicolon.

EXAMPLE:

```
operation execscript -sn SetKeyIntoGlobalParamTbl.sql -si DatabaseStateString -si y -si suspect;shutdown;standby
```

In this example, databases in the suspect, shutdown, and standby states will be automatically discovered.

FOR INDIVIDUAL CLIENTS

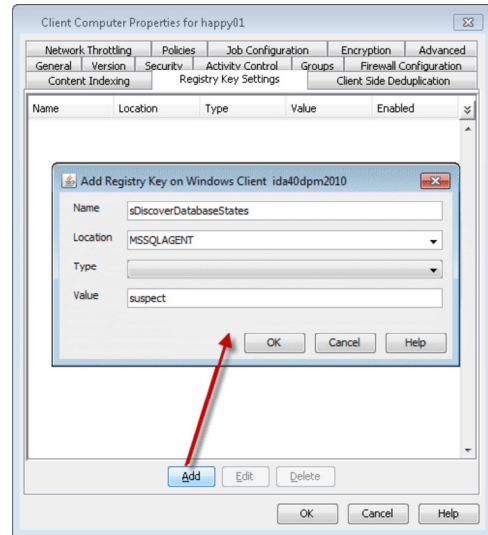
Configuring this option will override the configuration at the CommServe level described in the For All Clients section above.

1. From the CommCell Browser, 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 `sDiscoverDatabaseStates`.
6. In the **Location** list, type `MSSQLAgent`.
7. In the **Type** list, select `REG_MULTI_SZ`.
8. In the **Value** field, type the databases state or states that will be discovered. If entering more than one state, separate each with a semicolon.

For example, to discover databases in the Suspect, Shutdown, and Standby states, enter the following:

```
suspect;shutdown;standby
```

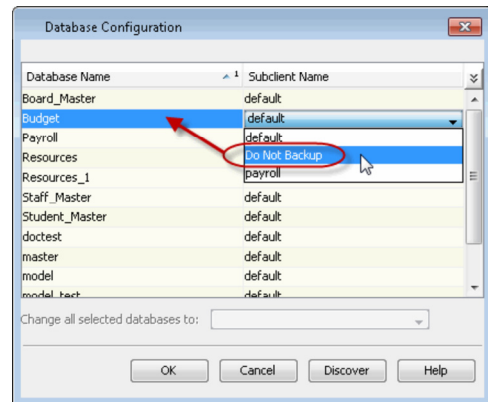
9. Click **OK**.



EXCLUDING DATABASES FROM BACKUPS

In some cases, it may be necessary to exclude certain databases from backups for a period of time. For example, you may have configured an entire SQL Server to back up using a particular schedule, but do not require all databases to be backed up according to that schedule. You can exclude databases from backups by following the steps below.

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server | <Instance>**.
2. Right-click the database subclient and click **Properties**.
3. Click the **Content** tab.
4. Click **Configure**.
5. Click **Discover**.
6. In the **Subclient Name** list, select the subclient to which the database you want to add should be assigned.
7. Click **OK** to save your settings.
8. Click **OK**.



MANAGING DATABASES DELETED FROM SQL SERVER

Databases that are deleted from the SQL Server are handled differently during backups depending on how the database was discovered and the type of subclient the database was originally assigned to.

Specifically:

- For default subclients:
 - Automatically discovered databases are automatically removed from the subclient upon deletion from the SQL Server.
 - Manually discovered databases must be manually removed from the subclient content.
- For user-defined subclients, automatically and manually discovered databases are not removed from the subclient upon deletion from the SQL Server. However, you can remove these databases manually from subclient content.

The following sections provide information on managing databases that have been deleted from the SQL Server.

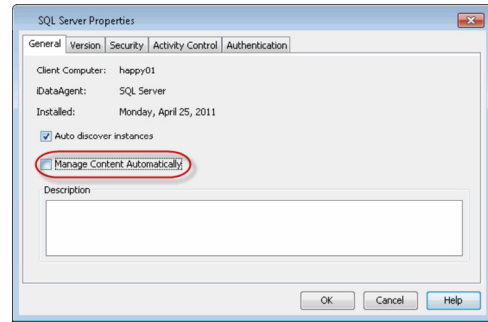
IN THE DEFAULT SUBCLIENT

AUTOMATICALLY DISCOVERED DATABASES

By default, when automatically discovered databases are deleted from the SQL Server, they are removed from the default subclient content during the subsequent backup.

In some cases, it may be desirable to manage all databases deleted from the SQL Server manually. For example, if a database is accidentally deleted from the SQL Server, it may not be desirable for it to be removed from the subclient during the next backup, especially if the next backup may occur while attempting to bring the deleted database back. Follow the steps below to manage deleted databases manually for the default subclient:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
2. Right-click **SQL Server** and click **Properties**.
3. Clear the **Manage Content Automatically** check box.
4. Click **OK**.

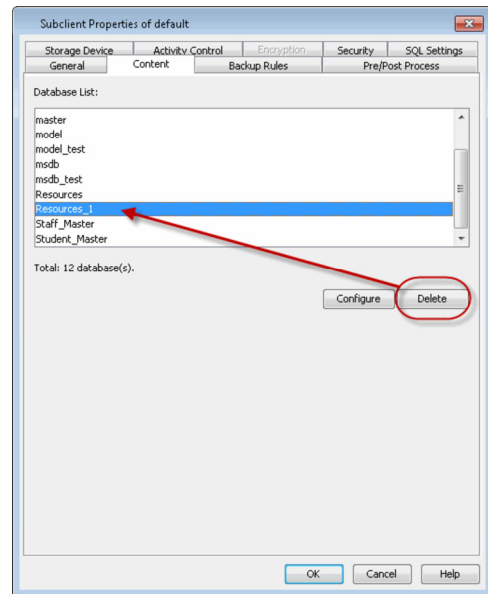


MANUALLY DISCOVERED DATABASES

Databases that have been manually added to the default subclient are not automatically removed from the subclient if deleted from the SQL Server. To ensure that the subsequent backup jobs complete without any errors, deleted databases from the server should be removed from the subclient.

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server | <Instance>**.
2. Right-click the default subclient and click **Properties**.
3. Click the **Content** tab.
4. Select the database you want to delete from the **Database List**.
5. Click **Delete**.
6. Click **OK**.

To automatically delete the databases from the subclient, set registry key `nIgnoreNonExistentDB` to 1.



IN USER-DEFINED AND DEFAULT SUBCLIENTS

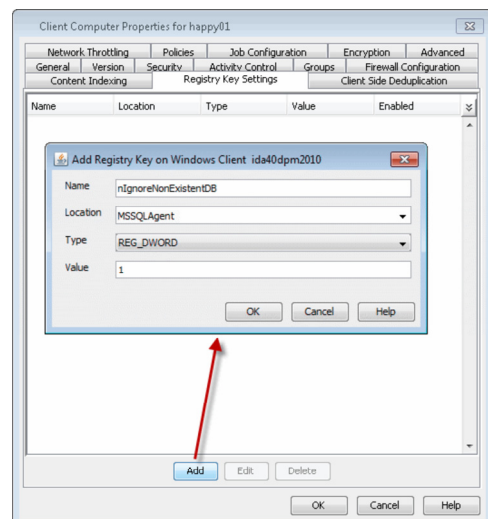
By default, databases are not automatically removed from subclients when deleted from the SQL Server. If a database is deleted from the SQL Server, the next backup performed on its associated subclient will not complete successfully.

You can also configure the subclient to automatically remove such databases as follows:

1. From the CommCell Browser, 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 `nIgnoreNonExistentDB`.

A global parameter `IgnoreNonExistentDB` can be used to configure all SQL clients. You can use the Command Line Interface to do so, see `qcommand execscript` for details.

6. In the **Location** list, type `MSSQLAGENT` from the dropdown box.
7. In the **Type** list, select **REG_DWORD**.
8. In the **Value** field, type 1.
9. Click **OK**.



SPECIFYING DATABASES FOR ONDEMAND BACKUP

OnDemand backup allows content to be specified as an external input at the time of initiating a data protection operation. Whereas traditional backups operations are performed on subclients, which have fixed content configured prior to performing the operation. On Demand backup allows you the flexibility of specifying content each time you perform a backup operation.

Use the following procedure to configure and perform OnDemand backup for Database subclients:

1. Download the CreateOnDemandDBSub.xml file and save it on the computer from where the command will be executed.
2. Execute the following command from the <Software_Installation_Directory>/Base folder after substituting the parameter values.

```
qoperation execute -af CreateOnDemandDBSub.xml -clientName client1 -instanceName client1\instance1 -subclientName
subclient1 -dataBackupStoragePolicy/storagePolicyName SPdata -logBackupStoragePolicy/storagePolicyName SPlog
```

where

clientName = client1,

instanceName = client1\instance1,

subclientName = subclient1,

dataStoragePolicy = SPData,

logStoragePolicy = SPLog.

3. Create a database list file (as a text file) for the databases you plan to backup. This database list file must be located on the Client computer, for which subclient needs to be created.

You may specify any databases to be backed up using the Database File List. There must be one entry per line in the file.

4. Download the backup_template.xml file and save it on the computer from where the command will be executed.
5. Execute the following command from the <Software_Installation_Directory>/Base folder after substituting the parameter values.

```
qoperation execute -af c:\backup_template.xml -clientName client1 -instanceName client1\instance1 -subclientName
subclient1 -ondemandinputfile c:\DatabaseFile.txt -backupLevel FULL
```

SELECTING BACKUP TYPES FOR ON DEMAND BACKUPS AND EXECUTING THE BACKUPS

You can modify the argument file (xml file) to perform different types of backups such as Full, Transaction Log and Differential backups, you can do so as follows:

1. The xml file will have the backup type parameter available, for example the Full backup parameter of the xml file (SQLbackup.bat):

```
<backupLevel>FULL</backupLevel>
```

can be changed to

```
<backupLevel>INCREMENTAL<backupLevel>
```

for Transaction Log backup

or to

```
<backupLevel>DIFFERENTIAL<backupLevel>
```

for Differential backup.

2. Execute the batch file using the following query in SQL Server Management Studio:

```
exec master..xp_cmdshell 'C:\SQLbackup.bat'
```

SPECIFYING FILE AND FILEGROUPS FOR ONDEMAND BACKUP

OnDemand backup allows content to be specified as an external input at the time of initiating a data protection operation. Whereas traditional backups operations are performed on subclients, which have fixed content configured prior to performing the operation. On Demand backup allows you the flexibility of specifying content each time you perform a backup operation.

Use the following procedure to configure and perform OnDemand backup for File and Filegroup Subclients:

1. Download the CreateOnDemandFFGSub.xml file and save it on the computer from where the command will be executed.
2. Execute the following command from the <Software_Installation_Directory>/Base folder after substituting the parameter values.

```
qoperation execute -af CreateOnDemandFFGSub.xml -clientName client1 -instanceName client1\instance1 -subclientName
subclient1 -dataBackupStoragePolicy/storagePolicyName SPdata -logBackupStoragePolicy/storagePolicyName SPlog
```

where

```

clientName = client1,
instanceName = client1\instance1,
subclientName = subclient1,
dataStoragePolicy = SPData,
logStoragePolicy = SPLog.
    
```

3. Create a file and file group list file (as a text file) for the files you plan to backup in the following format. Also ensure that the list file must be located on the Client computer, for which subclient needs to be created.

```

DatabaseName<\tab>FileGroupName<\tab>LogicalFileName
    
```

You may specify any files belonging to the same database only. There must be one entry per line in the file.

4. Download the backup_template.xml file and save it on the computer from where the command will be executed.
5. Execute the following command from the <Software_Installation_Directory>/Base folder after substituting the parameter values.

```

operation execute -af c:\backup_template.xml -clientName client1 -instanceName client1\instance1 -subclientName
subclient1 -ondemandinputfile c:\FileList.txt -backupLevel FULL
    
```

SETTING UP BACKUP CONVERSION RULES

Backup conversion rules provide the facility to convert certain types of backups to another backup type under specific circumstances. This functionality helps ensure all SQL data is protected regardless of circumstances that may cause a failure.

FOR DEFAULT AND DATABASE SUBCLIENTS

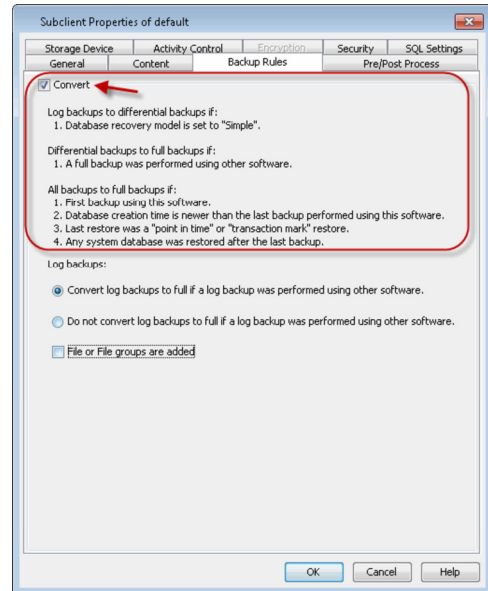
By default, database backups are converted as depicted in the following table.

BACKUP CONVERSION TYPE	CONDITIONS FOR CONVERSION	BENEFIT OF CONVERSION
Log Backup to Differential Backup	The database recovery model is set to Simple .	Because the Simple recovery model does not support log backups, converting to a differential backup ensures both logs and data are properly backed up. This, in turn, provides the facility to restore the logs.
Differential Backup to Full Backup	A full backups was performed using other software.	For first-time users, starting with a full backup provides complete protection as a baseline for future backups.
All Backups to Full Backups	<ul style="list-style-type: none"> • You are running your first backup using this software. • Database creation time is newer than the last backup performed using this software. • The last restore performed was a Point-in-Time or Transaction Mark restore. • Any system database (i.e., master, model, msdb) was restored after the last backup. 	Converting to full backups in these scenarios ensure you have complete protection of the latest state of each database. In the case of system databases, a full backup will ensure the restored database is backed up at the most recent point-in-time.

If you want to disable this functionality, you can do so using the steps below.

Keep in mind that disabling this option for one scenario disables the option for all scenarios listed above. As such, it is recommended this option remain enabled to ensure no data is unintentionally left out of a backup.

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server | <Instance>**.
2. Right-click the default subclient and click **Properties**.
3. Click the **Backup Rules** tab.
4. Disable the **Convert** check box.
5. Click **OK**.

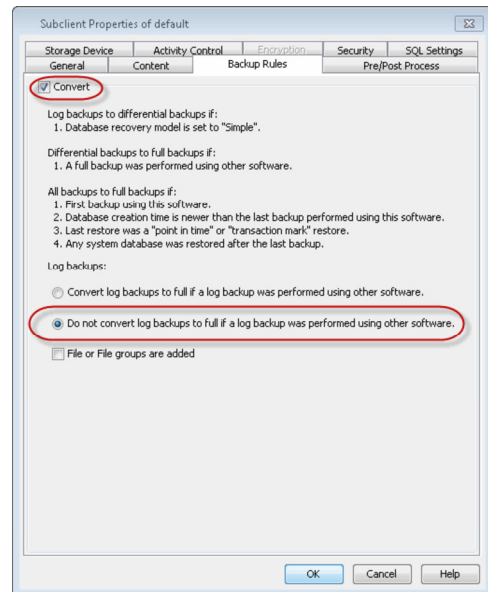


CONVERSION OPTIONS FOR LOG BACKUPS

By default, log backups performed outside of the system (for example, using SQL Enterprise Manager) are automatically converted to full backups. This provides a baseline for future backups.

If necessary, you can preserve the log backups performed by previous software packages as follows:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server | <Instance>**.
2. Right-click the default subclient and click **Properties**.
3. Click the **Backup Rules** tab.
4. Enable the **Convert** check box.
5. Select the **Do Not Convert Log Backups to Full if a Log Backup Was Performed Using Other Software** option.
6. Click the **SQL Settings** tab.
7. Select the **Disable Log Consistency Check** check box.
8. Click **OK**.



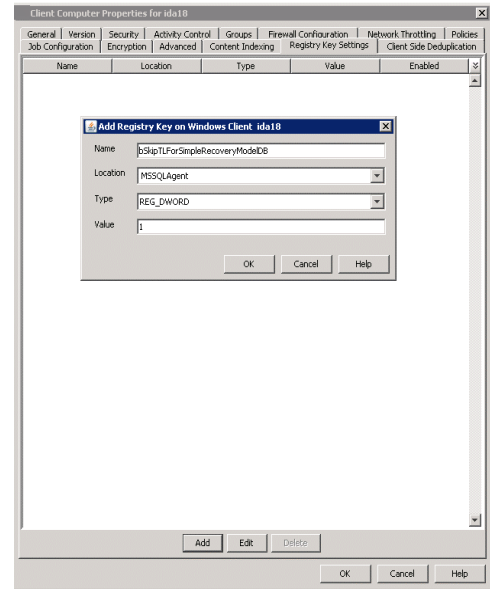
DISABLING CONVERSION OF TRANSACTION LOG BACKUPS TO DIFFERENTIAL

If the Convert checkbox is selected, all backups convert as specified in the rules of the dialog box. However, if you want to skip the conversion of Transaction Log Backups to differential backups for subclients with databases set to simple recovery model, you can do so by configuring additional settings to set bSkipTLForSimpleRecoveryModelDB.

1. From the CommCell Browser, navigate to **Client Computers**.
2. Right-click the **<Client>** in which you want to add the registry key, and then click **Properties**.
3. Click the **Registry Key Settings** tab.
4. Click **Add**.
5. Enter `bSkipTLForSimpleRecoveryModelDB` in the **Name** field.

A global parameter `SkipTLForSimpleRecoveryModelDB` can be used to configure all SQL clients. You can use the Command Line Interface to do so, see `qcommand` execscript for details.

6. Enter `MSSQLAgent` in the **Location** field.
7. From the **Type** list, select **REG_DWORD**.
8. In the **Value** field type `1`.
9. Click **OK**.

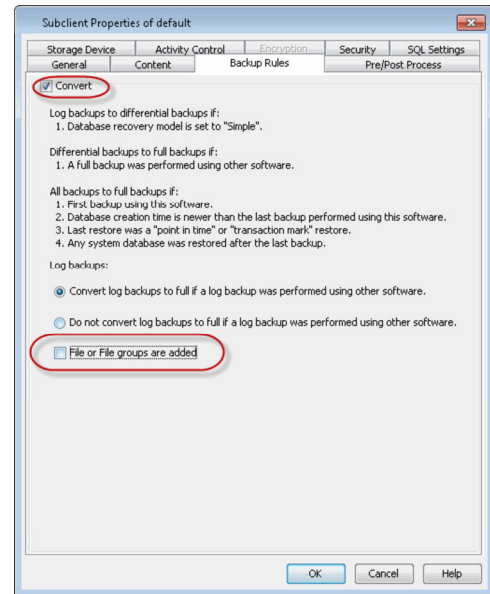


CONVERSION OPTIONS FOR FILES AND FILEGROUPS IN A DATABASE

By default, if files or filegroups have been added to a database since the previous backup, the next backup will automatically be converted to a full backup. This ensures the new files or filegroups are given proper protection as quickly as possible, regardless of the type of backup originally intended.

If you do not require this functionality, you can disable it as follows:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server | <Instance>**.
2. Right-click the default subclient and click **Properties**.
3. Click the **Backup Rules** tab.
4. Enable the **Convert** check box.
5. Clear the **File or Filegroups are added** check box.
6. Click **OK**.



FOR FILE/FILEGROUP SUBCLIENTS

By default, all backups performed on File/Filegroup subclients are automatically converted to full backups as depicted in the following table:

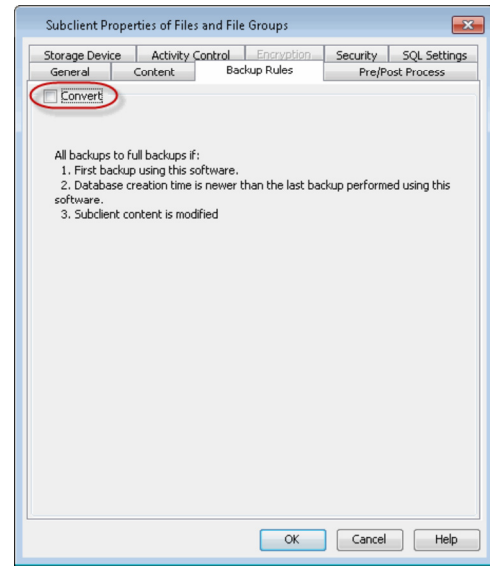
BACKUP CONVERSION TYPE	CONDITIONS FOR CONVERSION	BENEFIT OF CONVERSION
All Backups to Full Backups	<ul style="list-style-type: none"> • First backup using this software. • Database creation time is newer than the last backup performed using this software. • Subclient content is modified. 	Converting to full backups in these scenarios ensure you have complete protection of the latest state of each file/filegroup.

If you do not want backups to convert to full backups under these circumstances, you can disable this option by following the steps below.

Keep in mind that disabling this option for one scenario disables the option for all scenarios listed above. As such, it

is recommended this option remain enabled to ensure no data is unintentionally left out of a backup.

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server | <Instance>**.
2. Right-click the file or filegroup subclient and click **Properties**.
3. Click the **Backup Rules** tab.
4. Clear the **Convert** check box.
5. Click **OK**.



ENHANCING PERFORMANCE DURING BACKUPS

Several options are available for enhancing backup performance reducing network bandwidth overhead. These options include:

- Limiting the maximum size of data blocks used during backups.
- Specifying the number of buffers used to reserve bandwidth for data transfer.
- Limiting the maximum amount of data to be transferred at a time during backups.

You can configure these options as follows:

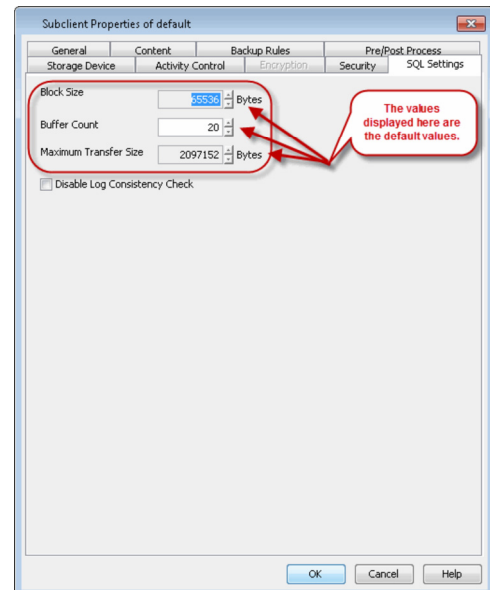
1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server | <Instance>**.
2. Right-click the default subclient and click **Properties**.
3. Click the **SQL Settings** tab.
4. Enter the desired number of data blocks to use during backups in the **Block Size** box.

All data transfers are in integral multiples of this value. The default value being 65,536 bytes (i.e., 64KB) or any value between 512 bytes and 65,536 bytes (inclusive) may be entered.

5. Enter the desired number of buffers to use during data transfer in the **Buffer Count** box.
The default value is 20.
6. Enter the maximum number of bytes to transfer at a time in the **Maximum Transfer Size** box.
The default value (in bytes) is 2097152.

- Make sure the Application Read Size value on the Data Transfer Option tab has the same or greater value as the Maximum Transfer Size; otherwise, backups may fail.
- The default value being 2,097,152 bytes (i.e., 2,048KB) or enter a value in multiples of 64 KB ranging between 65,536 bytes and 4,294,967,296 bytes (i.e., 4 MB)

7. Click **OK**.



CONFIGURING THE NUMBER OF LOG BACKUPS TO RUN BEFORE A FULL BACKUP

Full backups are necessary at regular intervals as it reduces the chance of data loss if one of log backup becomes corrupted as it will invalidate (not restorable) all other log backups performed after that. This key is used for the purpose of re-enforcing the need of a full backup after certain number of transaction log backups have run.

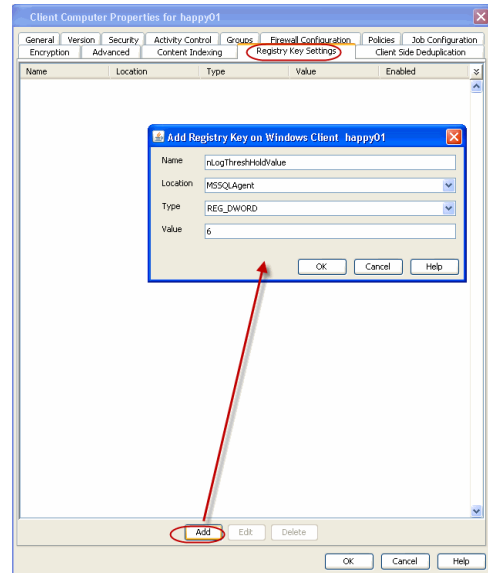
When this registry key is configured, a minor event will be generated in the Event Viewer to remind users to run a full backup after the configured number of transaction log backups have run.

Use the following steps to configure the number of log backups:

1. From the CommCell Browser, navigate to **Client Computers**.
2. Right-click the **<Client>** in which you want to add the registry key, and then click **Properties**.
3. Click the **Registry Key Settings** tab.
4. Click **Add**.
5. Enter `nLogThresholdValue` in the **Name** field.
6. Enter `MSSQLAgent` in the **Location** field.
7. From the **Type** list, select **REG_DWORD**.
8. Enter a number in the **Value** field. Range is [1 - <max_integer>].

This value specifies the number of transaction log backups that will be taken before a minor event is issued to remind users to run a full backup.

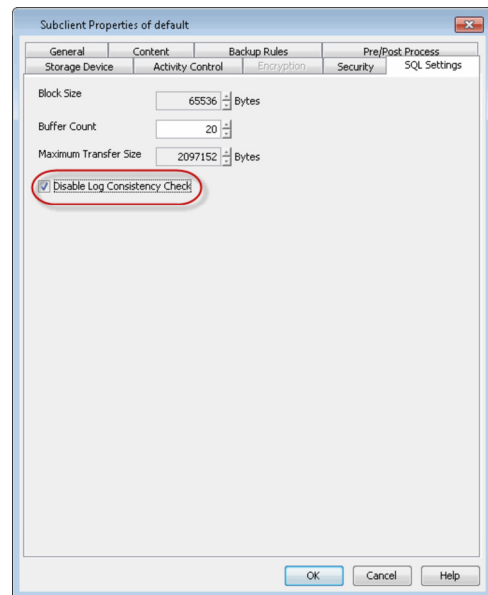
9. Click **OK**.



CONFIGURING LOG BACKUPS TO RUN WITHOUT FULL BACKUPS

By default, a full backup is required after performing a log backup. If you do not require a full backup at the time you want to back up the logs (for example, a full backup was performed outside of the system), you can do so as follows:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server | <Instance>**.
2. Right-click the subclient and click **Properties**.
3. Click the **SQL Settings** tab.
4. Enable the **Disable Log Consistency Check** check box.
5. Click **OK**.



PERFORMING BACKUPS USING VSS

In some cases, it may be desirable to allow write operations on SQL databases being backed up to continue while the backup runs. This is especially critical in environments where data entry to a SQL database is constant (such as databases being used daily in different parts of the world). Enabling VSS for backups provides this capability.

VSS can be enabled for backups of local volumes in both clustered and non-clustered environments. If the operating system fails to create a shadow copy of the data a traditional backup of the data will be performed, and a corresponding message will appear in the Event Viewer.

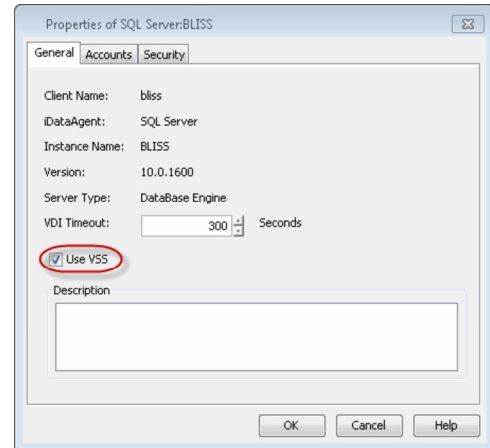
When VSS backups are performed, the following will automatically occur:

- All previously scheduled traditional full backups convert to VSS full backups.

- All previously scheduled differential and transaction log database backups are unaffected.
- All previously scheduled File/Filegroup backups are blocked from running.
- Immediate and scheduled VSS full backups run as single-stream backups.

During a VSS backup, the total amount of free space depends on the size of the backup data. As such, make sure to have sufficient disk space when you perform VSS backups.

1. Navigate to **Client Computers** | **<Client>** | **SQL Server** | **<Instance>**.
2. Right-click the instance and click **Properties**.
3. Enable the **Use VSS** check box.
4. Click **OK**.



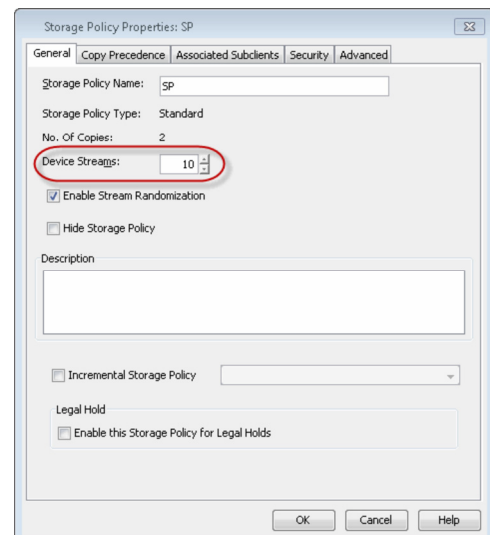
CONFIGURING DATA STREAMS

By default, backup data is sent to media in two streams. This means that a database, or a portion thereof, is sent to media during a backup in two parallel waves. This results in the backup taking about half the time to complete as it otherwise would if only one stream is used.

You can increase the number of streams used for backups for a particular subclient provided the number of streams does not exceed the maximum number configured in the subclient's storage policy. Increasing the number of streams for a subclient further reduces the amount of time a backup takes to complete. For example, increasing the number of streams from 2 to 3 enhances backup time from one-half that of a single stream to one-third.

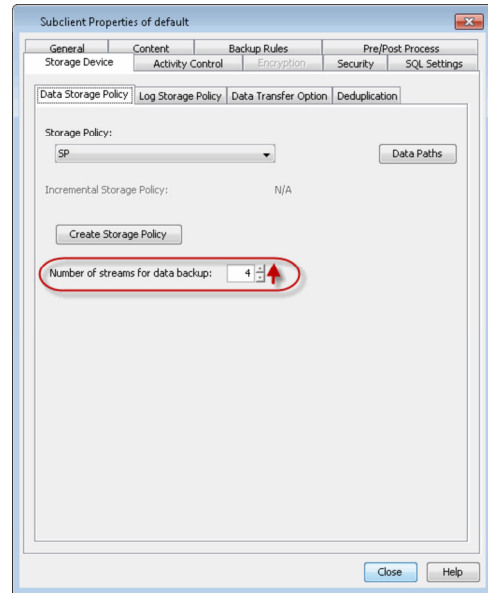
Keep in mind that the number of streams configured for backups must also be used when restoring data. For example, if you configure a subclient to use 4 streams, you must also use 4 streams to restore the data.

1. From the CommCell Browser, navigate to **Policies** | **Storage Policies**.
2. Right-click the storage policy associated with the subclient you want to increase the streams for and click **Properties**.
3. Ensure the number in the **Device Streams** box is greater than the number of streams you want to configure for the subclient.



1. Navigate to **Client Computers** | **<Client>** | **SQL Server** | **<Instance>**.
2. Right-click the subclient and click **Properties**.
3. Click the **Storage Device Tab** tab.
4. Increase (or decrease) the number of streams in the **Number of Streams for data backup** box.
5. Click the **Log Storage Policy** tab.

6. Increase (or decrease) the number of streams in the **Number of Streams for transaction log** box.
7. Click **OK**.



CONFIGURING USER ACCOUNTS FOR BACKUPS

The SQL Server iDataAgent requires a Windows user account that has sufficient privileges for the software to:

- Perform backups and restores
- Access the Windows registry
- Stop or start the SQL Server services.

The following table illustrates the requirements for the user account you will need to provide:

IF THE SQL SERVER IS:	THE USER ACCOUNT SHOULD BE:
On a non-Domain Controller	<ul style="list-style-type: none"> • Local Administrator of the computer on which the SQL Server resides. • Member of the SQL sysadmin fixed server role. • User account from which SQL Server services are running. If this account differs from the Local Administrator account, you must provide the credentials of the account used by the SQL Server services.
On a Domain Controller	An account other than the Domain Administrator account that has Administrator and SQL sa privileges.

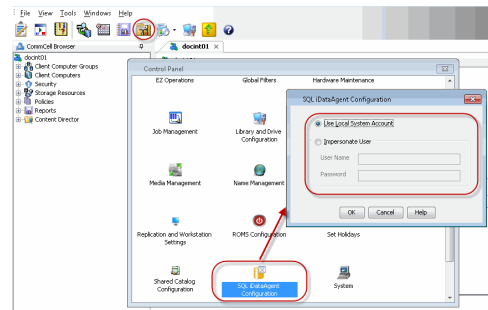
By default, the local system administrator account is used for all components. You can change the user account to be used for backups at the CommCell, Client Computer Group, Agent, and Instance levels. Accounts configured at each level will be used for all entities within that level as described in the following sections.

In order to access the SQL Server databases to perform data protection and recovery operations, the SQL **sysadmin** rights are required.

AT THE COMMCELL LEVEL

This user account will be used for all SQL Server iDataAgents in your CommCell. Configure the user account at this level if one person will be conducting all backup and restore operations in your organization.

1. Navigate to **Control Panel**.
2. Double-click the **SQL iDataAgent Configuration** icon.
3. Select the following:
 - Use Local System Account** if the computer's Administrator account contains the required privileges.
 - Impersonate User** if you want to use a different account that contains the required privileges. Enter the **User Name** and **Password** for this account in the space provided.
4. Click **OK**.

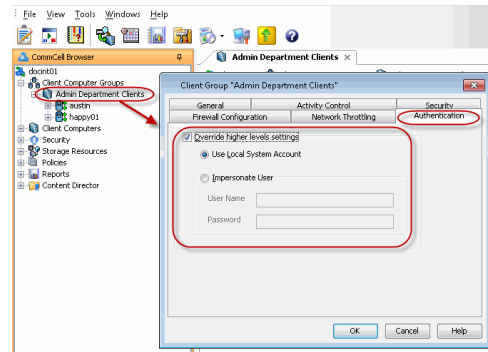


AT THE CLIENT COMPUTER GROUP LEVEL

This user account will be used for all computers within a Client Computer Group. Configure the user account at this level if different people will be conducting

backup and restore operations for each Client Computer Group in your organization. This user account will override the user account configured at the CommCell level.

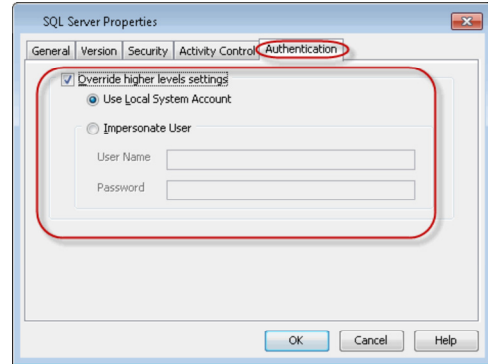
1. Navigate to **Client Computer Groups**.
2. Verify that all the SQL Server clients for which you wish to configure the user account is included in the Client Computer Groups.
3. Right-click the Client Computer Group and click **Properties**.
4. Enable the **Override higher levels settings** check box.
5. Select the following:
 - Use Local System Account** if the computer's Administrator account contains the required privileges.
 - Impersonate User** if you want to use a different account that contains the required privileges. Enter the **User Name** and **Password** for this account in the space provided.
6. Click **OK**.



AT THE AGENT LEVEL

This user account will be used for all instances and associated subclients. Configure the user account at this level if one person will be conducting all backup and restore operations on the client on which the SQL Server iDataAgent is installed. This user account will override the user account configured at the CommCell and Client Computer Group levels.

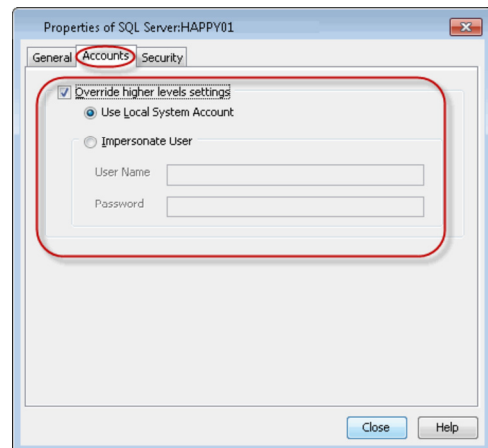
1. Navigate to **Client Computers | <Client>**.
2. Right-click **SQL Server** and click **Properties**.
3. Click the **Authentication** tab.
4. Enable the **Override higher levels settings** check box.
5. Select the following:
 - Use Local System Account** if the computer's Administrator account contains the required privileges.
 - Impersonate User** if you want to use a different account that contains the required privileges. Enter the **User Name** and **Password** for this account in the space provided.
6. Click **OK**.



AT THE INSTANCE LEVEL

This user account will be used for all subclients within the instance. Configure the user account at this level if backup and restore operations will be conducted by a different person for each instance. This user account will override the user account configured at the CommCell, Client Computer Group, and Agent levels.

1. Navigate to **Client Computers | <Client> | SQL Server | <Instance>**.
2. Right-click the instance and click **Properties**.
3. Click the **Accounts** tab.
4. Enable the **Override higher levels settings** check box.
5. Select the following:
 - Use Local System Account** if the computer's Administrator account contains the required privileges.
 - Impersonate User** if you want to use a different account that contains the required privileges. Enter the **User Name** and **Password** for this account in the space provided.
6. Click **OK**.



COMMAND LINE OPERATIONS

You can add, modify or delete several configurable properties for SQL iDataAgent from command line.

Command line configuration enables you to:

- configure the same properties across multiple clients simultaneously.
- reuse the same configurations for additional entities.

The following sections describe the available command line configurations:

LOG ON TO THE COMMSERVE

To run command line operations you must first login to the CommServe as follows:

- From Command prompt, navigate to <Software_Installation_Directory>/Base and run the following command:

```
qlogin -cs <commserve name> -u <user name>
```

- For example, to log on to CommServe 'server1' with username 'user1':

```
qlogin -cs server1 -u user1
```

CONFIGURE INSTANCES

CREATE AN INSTANCE

1. Download the CreateSQL_Instance_Template.xml file and save it on the computer from where the command will be executed.
2. Execute the following command from the <Software_Installation_Directory>/Base folder after substituting the parameter values.

```
qoperation execute -af CreateSQL_Instance_Template.xml -appName 'SQL Server' -clientName xxxxx -instanceName xxxxx
```

MODIFY AN INSTANCE

1. Download the ModifySQL_Instance_Template.xml file and save it on the computer from where the command will be executed.
2. Execute the following command from the <Software_Installation_Directory>/Base folder after substituting the parameter values.

```
qoperation execute -af ModifySQL_Instance_Template.xml -appName 'SQL Server' -clientName xxxxx -instanceName xxxxx
```

GET INSTANCE PROPERTIES

1. Download the GetSQL_Instance_Properties_Template.xml file and save it on the computer from where the command will be executed.
2. Execute the following command from the <Software_Installation_Directory>/Base folder after substituting the parameter values.

```
qoperation execute -af GetSQL_Instance_Properties_Template.xml -appName 'SQL Server' -clientName xxxxx -instanceName xxxxx
```

DELETE AN INSTANCE

1. Download the DeleteSQL_Instance_Template.xml file and save it on the computer from where the command will be executed.
2. Execute the following command from the <Software_Installation_Directory>/Base folder after substituting the parameter values.

```
qoperation execute -af DeleteSQL_Instance_Template.xml -appName 'SQL Server' -clientName xxxxx -instanceName xxxxx
```

AVAILABLE COMMAND PARAMETERS FOR INSTANCES

The following table displays all the parameters that you can use with the commands mentioned in the above sections. To add a parameter to your command, use the following syntax: (An example is provided at the end of the table.)

```
qoperation execute -af <template XML file> -<parameter name> <value>
```

PARAMETER	DESCRIPTION OF PARAMETER VALUES
description	A general description of the instance if needed.
clientName	Name of the client computer.
appName	Name of the application. In this case it would be 'SQL Server'.
instanceName	Name of the SQL instance. If you are specifying an instance other than the default instance, the name must be entered along with the client name, e.g. client1\instance1.
vDITimeOut	The time the system must wait for the SQL server to become ready. Time value is in seconds.
useVss	Option to enable/disable VSS. Valid values are True/False.
overrideGlobalAuthentication	Option to override Global Authentication. Valid values are True/False.
useLocalSystemAccount	Option to use Local Administrator account. Valid values are 0/1.
userName	Name of the database administrator who has permissions to perform jobs.
password	Password for the database user
propertyLevel	Properties to be retrieved for the instance.

Valid values are:

- ListOnly - List the instance names
- BasicProperties - Retrieves the basic properties, such as description, version, etc.,
- ExtendedProperties - Retrieves all the properties for the instance.
- AllProperties - Retrieves all the properties for the instance.

The following example shows how to add a parameter for a command:

Enable VSS	To enable VSS for the instance, add the 'useVSS' parameter to the following command: <pre>goperation execute -af ModifySQL_Instance_Template.xml -appName 'SQL Server' -clientName client1 -instanceName client1\instance1 -useVSS true</pre>
-------------------	--

For detailed information on the instance parameters see Instance Properties.

CONFIGURE SUBCLIENTS

CREATE SUBCLIENT

1. Download the CreateSQL_Subclient_Template.xml file and save it on the computer from where the command will be executed.
2. Execute the following command from the <Software_Installation_Directory>/Base folder after substituting the parameter values.

```
goperation execute -af CreateSQL_Subclient_Template.xml -appName 'SQL Server' -clientName xxxxx -instanceName xxxxx -subclientName xxxxx -dataBackupStoragePolicy/storagePolicyName xxxxx -logBackupStoragePolicy/storagePolicyName xxxxx
```

GET SUBCLIENT PROPERTIES

1. Download the GetSQL_Subclient_Properties_Template.xml file and save it on the computer from where the command will be executed.
2. Execute the following command from the <Software_Installation_Directory>/Base folder after substituting the parameter values.

```
goperation execute -af GetSQL_Subclient_Properties_Template.xml -appName 'SQL Server' -clientName xxxxx -instanceName xxxxx -subclientName xxxxx
```

MODIFY SUBCLIENT

1. Download the ModifySQL_Subclient_Template.xml file and save it on the computer from where the command will be executed.
2. Execute the following command from the <Software_Installation_Directory>/Base folder after substituting the parameter values.

```
goperation execute -af ModifySQL_Subclient_Template.xml -appName 'SQL Server' -clientName xxxxx -instanceName xxxxx -subclientName xxxxx
```

DELETE SUBCLIENT

1. Download the remove_subclient_template.xml file and save it on the computer from where the command will be executed.
2. Execute the following command from the <Software_Installation_Directory>/Base folder after substituting the parameter values.

```
goperation execute -af remove_subclient_template.xml -appName 'SQL Server' -clientName xxxxx -instanceName xxxxx -subclientName xxxxx
```

AVAILABLE COMMAND PARAMETERS FOR SUBCLIENTS

The following table displays all the parameters you can use with the commands mentioned in the above sections. To add a parameter to your command, use the following syntax: (Some examples are provided at the end of the table.)

```
goperation execute -af <template XML file> -<parameter name> <value>
```

PARAMETER	DESCRIPTION OF PARAMETER VALUES
appName	Name of the application. In this case it would be 'SQL Server'.
backupsetName	Name of the backupset. In this case, it would be 'defaultBackupSet'.
clientName	Name of the client computer.
subclientName	Name of the SQL subclient.
instanceName	Name of the SQL instance. If you are specifying an instance other than the default instance, the name must be entered along with the client name, e.g. client1\instance1.
enableBackup	Option to enable backup of Subclient. Valid values are True/False.
encryptionFlag	Option to set the encryption points during backups. Valid values are: <ul style="list-style-type: none"> • ENC_MEDIA_ONLY, to encrypt the backup data after transmission and prior to storage on the media. • ENC_NETWORK_AND_MEDIA, to encrypt the backup data before transmission. The data is stored encrypted on the media. • ENC_NETWORK_ONLY, to encrypt the backup data for transmission and then decrypt the data prior to storage on the media.

	<ul style="list-style-type: none"> ENC_NONE, to disable data encryption.
readBuffersize	<p>This is used for tuning SQL Server performance.</p> <p>Valid values are 64, 128, 256, 512, 1024, 2048 and 4096.</p>
onDemandSubClient	<p>Use this option during creation of a subclient.</p> <p>Valid values are True/False.</p>
networkAgents	Number of Network Agents.
softwareCompression	<p>Option to enable compression on the Client or MediaAgent computer. Valid values are:</p> <ul style="list-style-type: none"> ON_CLIENT, to enable software compression on the client. ON_MEDIAAGENT, to enable software compression on the MediaAgent. OFF, to disable software compression.
throttleNetworkBandwidth	<p>Enhancing backup performance by reducing network bandwidth overhead.</p> <p>Valid values are 0/1.</p>
storagePolicyName	Name of Storage Policy to be associated for Backup.
enableDeduplication	<p>Option to enable deduplication on subclient.</p> <p>Valid values are True/False.</p>
generateSignature	<p>A component of deduplication performed on the client or MediaAgent computer. Valid values are:</p> <ul style="list-style-type: none"> ON_CLIENT, to enable signature generation on the client. ON_MEDIA_AGENT, to enable signature generation on the MediaAgent. OFF, to disable signature generation.
backupRules	<p>Option to set Backup conversion rules. Valid values are:</p> <ul style="list-style-type: none"> CONVERT_WITH_FILE_FILEGROUPS_NOT_ADDED, to convert log backups to full if a log backup was performed using other software with file or file group not added. CONVERT_WITH_FILE_FILEGROUPS_ADDED, to convert log backups to full if a log backup was performed using other software with file or file group added. DONOT_CONVERT_WITH_FILE_FILEGROUPS_NOT_ADDED, do not convert log backups to full if a log backup was performed using other software with file or file groups are not added. DONOT_CONVERT_WITH_FILE_FILEGROUPS_ADDED, do not convert log backups to full if a log backup was performed using other software with file or file groups are added. CONVERT_DISABLED, to disable backup conversion rules.
bufferCount	Number of data blocks to use during backups to improve performance.
maxTransferSize	<p>Maximum number of bytes to transfer at a time to tune performance.</p> <p>Valid values in bytes are 65536, 131072, 262144, 524288, 1048576, 2097152, and 4194304.</p>
numberOfBackupStreams	Number of Backup Streams.
numberOfTransactionLogStreams	Number of Transaction Log Streams.
runPostBackup	<p>Option to run a process after backup completes.</p> <p>Valid values are Yes/No.</p>
postBackupCommand	Path to the post process script that will be run after the backup.
preBackupCommand	Path to the pre process script that will be run before the backup.
runAs	<p>Option to specify the user name who has permissions to run the pre/post process scripts. Valid values are:</p> <ul style="list-style-type: none"> USE_IMPERSONATION, to specify a user with enough privileges to run the scripts. When using this value, you also need to provide the user credentials: -prepostUserName/userName <user name> -prepostUserName/password <password> USE_LOCAL_SYS_ADMIN, to use the administrator account to run the scripts.
sqlSubclientType	<p>Type of SQL subclient. Valid values are:</p> <ul style="list-style-type: none"> DATABASE, to back up groups of specific databases. FILE_FILEGROUP, to back up portions of data from large databases.
contentOperationType	<p>Modification type. Valid values are</p> <ul style="list-style-type: none"> ADD, to add content to the subclient OVERWRITE, to overwrite contents to the subclient. DELETE, to delete the contents from the subclient.
disableLogConsistencyCheck	<p>Option to disable log consistency check.</p> <p>Valid values are True/False.</p>

The following examples show how to add an parameter for a command:

Enable Backup	<p>To enable backup for a subclient, add the 'enableBackup' parameter to the following command:</p> <pre>qoperation execute -af ModifySQL_Subclient_Template.xml -appName 'SQL Server' -clientName client1 -instanceName client1\instance1 -subclientName subclient1 -enableBackup true</pre>
Setting a Storage Policy	<p>To assign a storage policy to a subclient, add the 'storagePolicyName' parameter to the following command:</p> <pre>qoperation execute -af modify_storage_policy.xml -appName 'SQL Server' -clientName client1 -instanceName client1\instance1 -subclientName subclient1 -dataBackupStoragePolicy/storagePolicyName SPdata -logBackupStoragePolicy/storagePolicyName SPlog</pre>
Creating an On Demand	<p>To create an on demand database subclient, add the 'onDemandSubClient' and 'sqlSubclientType' parameters to the following command:</p>

Database Subclient	<pre>operation execute -af CreateSQL_Subclient_Template.xml -appName 'SQL Server' -clientName client1 -instanceName client1\instance1 -subclientName subclient1 -onDemandSubClient true -sqlSubclientType DATABASE -dataBackupStoragePolicy/storagePolicyName SPData -logBackupStoragePolicy/storagePolicyName SPLog</pre>
Creating an On Demand File/FileGroup Subclient	<p>To create an on demand file/filegroup subclient, add the 'onDemandSubClient' and 'sqlSubclientType' parameters to the following command:</p> <pre>operation execute -af CreateSQL_Subclient_Template.xml -appName 'SQL Server' -clientName client1 -instanceName client1\instance1 -subclientName subclient1 -onDemandSubClient true -sqlSubclientType FILE_FILEGROUP -dataBackupStoragePolicy/storagePolicyName SPData -logBackupStoragePolicy/storagePolicyName SPLog</pre>
Adding Content to a Subclient	<p>To add content to a subclient, add the 'contentOperationType' parameter to the following command:</p> <pre>operation execute -af modify_subclient_content.xml -appName 'SQL Server' -clientName client1 -instanceName client1\instance1 -subclientName subclient1 -sqlSubclientType DATABASE -mssqlDbContent/databaseName DB1 -contentOperationType ADD</pre>
Overwriting Content to a Subclient	<p>To overwrite content to a subclient, add the 'contentOperationType' parameter to the following command:</p> <pre>operation execute -af modify_subclient_content.xml -appName 'SQL Server' -clientName client1 -instanceName client1\instance1 -subclientName subclient1 -sqlSubClientType DATABASE -mssqlDbContent/databaseName DB1 -contentOperationType OVERWRITE</pre>

For detailed information on the Subclient parameters see Subclient Properties.

PERFORM COMMAND LINE OPERATIONS FROM MICROSOFT SQL MANAGEMENT STUDIO

Besides using the operating system's command line interface, you can also perform the command line operations from MSSQL Management Studio interface by integrating the commands within the SQL scripts. Registration of stored procedure and execution of commands should be done from SQL Server version 2008 or higher.

Use the following steps to run command line operations from SQL Management Studio:

1. Execute the query to register the stored procedure to a specific SQL database (eg., msdb) in MSSQL Management Studio.

In order to do this, make sure the trustworthy database property for the SQL database is set to ON.

```
DECLARE @i_FullDLLPath NVARCHAR(MAX)
SET @i_FullDLLPath = 'C:\Calypso\Base\dbclr.dll'--INPUT REQUERED. Enter
dbclr.ll PATH. eg: <software_install_path>\Base\dbclr.dll
DECLARE @dbName nvarchar(255)
SET @dbName ='msdb'--INPUT REQUERED. Give DATABASAE name where this
procedure is needed.

DECLARE @o_retVal INTEGER;
```

2. Log in to the CommServe using the SQL scripts. For example,

```
USE [msdb]
GO
DECLARE @return_value int,
@response nvarchar(max)
EXEC @return_value = [dbo].[cv_cmdshell]
@input = N'qlogin -u "commandline" -
ps "3d4b14baf641d429e2b3782c1ed5d0a64" -cs "commserve.companyname.com",
@response = @response OUTPUT
```

You can login to the CommServe using an encrypted or open password.

- o For encrypted password, run a save as script operation from the CommCell Console and view the .bat file to retrieve the encrypted password.
- o Open password can be provided as shown below:

```
@input = N'qlogin -u "admin" -clp "admin",
```

3. Execute the command line operation using SQL scripts.

Examples:

Creating an Instance	<pre>DECLARE @return_value int, @response nvarchar(max) EXEC @return_value = [dbo].[cv_cmdshell] @input = N'qlist instance -c client1 -a Q_MSSQL', @response = @response OUTPUT IF RTRIM(@response) LIKE N'client1\instance1%'</pre>
-----------------------------	--

	<pre>Print 'Instance already exists' ELSE Print 'Creating instance client1\instance1' EXEC @return_value = [dbo].[cv_cmdshell] @input = N'qoperation execute -af e:\MyFolder\createSQL_Template.xml -instanceName client1\instance1 -clientName client1', @response = @response OUTPUT where, e:\MyFolder is the location where the template XML is saved.</pre>
<p>Performing a full backup</p>	<pre>DECLARE @return_value int, @response nvarchar(max) EXEC @return_value = [dbo].[cv_cmdshell] @input = N'qoperation execute -af e:\MyFolder\full.xml -subClientName subclient1 -clientName client1 -instanceName client1\instance1 -ondemandinputfile C:\MyTest\myDBsContent.txt', @response = @response OUTPUT where, myDBsContent.txt is the content file that list the databases to be backed up and e:\MyFolder is the location where the template XML is saved.</pre>
<p>Perform a restore</p>	<pre>DECLARE @return_value int, @response nvarchar(max) EXEC @return_value = [dbo].[cv_cmdshell] @input = N'qoperation execute -af e:\MyFolder\restore_template.xml -clientName client1 -instanceName client1\instance1 -toTimeValue 2011-11-28 15:40:00 -restoreSource DB1 -database DB1', @response = @response OUTPUT where e:\MyFolder is the location where the template XML is saved and the -toTimeValue is the backup finish date.</pre>

UN-REGISTERING THE SQL DATABASE

You can un-register the stored procedure from the SQL database by executing the below query:

```
DECLARE @dbName nvarchar(255)
SET @dbName = 'msdb'--INPUT REQUERED. Give DATABASEAE name where this procedure is needed.
DECLARE @dropAssembly int = 1 -- 1 = true , 0 = false

DECLARE @o_retVal INTEGER;

DECLARE @sqlQuery NVARCHAR(MAX)
```

MODIFYING AN AGENT, INSTANCE, OR SUBCLIENT

There are several configurable properties available for your agent that can be modified from the agent, instance, or subclient level as per need.

It is recommended that you do not modify the properties of a subclient when a job is in progress for that specific subclient. If a job is in progress, either wait for the job to complete or kill the job from the Job Controller.

The following table describes the properties that can configured from the agent, instance, and subclient levels.

OPTION	DESCRIPTION	RELATED TOPICS
<p>Change Storage Policies</p>	<p>You can modify the storage policies in any of the following situations:</p> <ul style="list-style-type: none"> To include a different media for the backup operation. To use a storage policy with a different retention criteria. <p>You can change the storage policies from the subclient level.</p> <ol style="list-style-type: none"> From the CommCell Browser, right-click the subclient. Click Properties. Click Storage Device. Select the Storage policy from the drop-down menu. 	<p>Refer to Storage Policies.</p>

	5. Click OK .	
Rename a Subclient	<p>You can rename a subclient:</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, right-click the subclient. 2. Click Properties. 3. Type the new name in the Subclient name field. 4. Click OK. 	
Data Transfer Options	<p>You can efficiently configure the available resources for transferring data secured by data protection operations from the subclient level. This includes the following:</p> <ul style="list-style-type: none"> • 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. <p>You can configure the data transfer options.</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, right-click the subclient. 2. Click Properties. 3. Click Storage Device. 4. Click Data Transfer Option tab. 5. Choose the appropriate software compression option for this subclient. 6. Select Throttle Network Bandwidth and set the required bandwidth. 7. Click OK. 	Refer to Data Compression and Network Bandwidth Throttling.
View Data Paths	<p>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 including their priority from the subclient level.</p> <ol style="list-style-type: none"> 1. From the CommCell browser, right-click the subclient. 2. Click Properties. 3. Click Storage Device. 4. Select Storage Policy from the drop-down menu. 5. Click Data Paths. 	
Configure a Subclient for Pre/Post Processing of Data Protection	<p>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.</p> <ol style="list-style-type: none"> 1. From the CommCell browser, right-click the subclient. 2. Click Properties. 3. Click Pre/Post Process. 4. Click one of the following phases and type the full path of the process that you want to execute during that phase. Alternatively, click Browse to locate the process (applicable only for paths that do not contain any spaces). <ul style="list-style-type: none"> o PreBackup Process o PostBackup Process 5. Click OK. 6. Select Run Post Backup Process for all attempts to run a post backup process for all attempts. 7. For subclients on Windows platforms, Run As displays Not Selected. If you want to change the account that has permission to run these commands, click Change. <ol style="list-style-type: none"> a. In the User Account dialog box, select Use Local System Account, or select Impersonate User and enter the 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. 	Refer to Pre/Post Processes.
Configure Activity Control	<p>You can enable backup and restore operations from the agent and subclient level. However, you can enable restore operations only from the agent level.</p> <ol style="list-style-type: none"> 1. From the CommCell browser, right-click the subclient. 2. Click Properties. 3. Click Activity Control, select or clear option(s) as desired. 4. Click OK. 	Refer to Activity Control.
Configure User Security	<p>You can configure user security from the agent or subclient level.</p> <p>You can perform the following functions:</p> <ul style="list-style-type: none"> • Identify the user groups to which this CommCell object is associated. 	Refer to User Administration and Security.

	<ul style="list-style-type: none"> • Associate this object with a user group. • Disassociate this object from a user group. <ol style="list-style-type: none"> 1. From the CommCell browser, right-click the subclient. 2. Click Properties. 3. Click Security. 4. Select the appropriate user groups to which you want to associate to the CommCell object from the Available Groups pane, and then move the user group to the Associated Groups pane. 5. Click OK. 	
Enable/Disable Data Encryption	<p>You can enable data encryption from the subclient level. Encryption must be enabled at the client level prior to configuring any instances residing on that client.</p> <ol style="list-style-type: none"> 1. From the CommCell browser, right-click the subclient. 2. Click Properties. 3. Click Encryption. 4. Select the desired encryption. 5. Click OK. 	Refer to Data Encryption.
View Software Version and Installed Updates	<p>The Version tab, at the Agent level displays the software version of the component.</p> <ol style="list-style-type: none"> 1. From the CommCell browser, right-click the agent. 2. Click Properties. 3. Click Version. 4. Click OK. 	
CommCell Configuration Report	<p>The CommCell Configuration Report provides the properties of the CommServe, MediaAgents, clients, agents, SRM agents, subclients, and storage policies within the CommCell based on the selected filter criteria.</p> <ol style="list-style-type: none"> 1. From the CommCell browser, click Reports icon. 2. Select CommCell Configuration. 3. Click Run. 	Refer to CommCell Configuration.

DELETING AN AGENT, INSTANCE, OR SUBCLIENT

The following sections describe the steps involved in deleting an agent, instance, or subclient.

When you delete an instance or backupset, the associated data is logically deleted and you can no longer access the corresponding data from CommCell Console for recovery purposes.

Refer to the troubleshooting article on Recovering Data Associated with Deleted Clients and Storage Policies for information on how to recover data if you accidentally delete an entity.

DELETING AN AGENT

You need to uninstall or DeConfigure the agent software from the client computer before deleting from CommCell Browser. After you delete the client software, you can either leave the corresponding data intact for appropriate action or you can remove the data immediately. If you choose to remove the data immediately, you must delete the agent from the CommCell Browser. If you delete the agent, all of the agent's data is irretrievably lost.

- You cannot delete an agent while operations for that agent are running.
1. From the CommCell Browser, navigate to **Client Computers | <Client>**.
 2. Right-click the **<Agent>**, and then click **Delete**.

3. A confirmation message is displayed with the following message:

```
This operation will permanently delete the data backed up from this level and it cannot be restored.
```

4. Click **OK** to continue with the deletion operation., or click **No** to abort the deletion.

DELETING AN INSTANCE

Consider the following before deleting an instance:

- When you delete a specific instance all job schedules and job histories that pertain to any of the levels within the deleted instance are deleted.
- You cannot delete an instance if it is being backed up. Attempts to delete an instance under such conditions cause the deletion to fail. If a backup is in progress, either wait for the backup to complete or kill the backup job using the Job Manager. Once the backup is no longer in progress, you can delete the instance level.
- You cannot delete an instance if there is only one instance present for an agent. To delete the final instance, you must remove the agent software from the

client computer.

1. From the CommCell Browser, right-click the instance that you want to delete, click **All Tasks** and then click **Delete**.
2. Click **Yes** to confirm the deletion. (Clicking **No** cancels the deletion and retains the node.)
3. Type the requested phrase in the **Enter Confirmation Text** dialog box and click **OK**. This should delete the instance.

DELETING A SUBCLIENT

Consider the following before deleting a subclient:

- You cannot delete a default subclient.
 - Schedules associated with the subclient are also automatically deleted.
1. From the CommCell Browser, navigate to **Client Computers | <Client> | <Agent> | <Instance>**
 2. Right-click the **<subclient>** that you want to delete, and then click **Delete**.
 3. 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.

Advanced Backup – SQL Server iDataAgent

TABLE OF CONTENTS

<p>Full Backups</p> <p>Transaction Log Backups</p> <ul style="list-style-type: none"> Automatically Start a Transaction Log Backup after a Successful Backup Run Transaction Log Backups Without Running a Full Backup Back Up Transaction Log of a Damaged Database (Do Not Truncate Logs) Tail Log Backup Configure Number of Log Backups Before Running Full <p>Differential Backups</p> <p>Compressing Backups with Native SQL Compression (for SQL Server 2008 or later)</p> <p>Perform Partial Backup to Exclude Read-Only FileGroups</p> <p>Scheduling a Backup</p> <p>Managing Jobs</p> <ul style="list-style-type: none"> Restarting Jobs Controlling Jobs <p>Additional Options</p>	<p>Command Line Backups</p> <ul style="list-style-type: none"> Log on to the CommServe Perform the Backup Generate the Command Line Script from CommCell Console
--	--

FULL BACKUPS

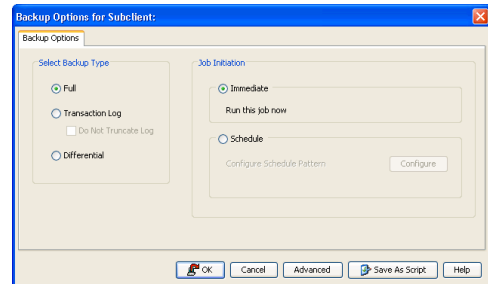
Full backups provide the most comprehensive protection of data.

Backups for any client start with a full backup. The full backup becomes a baseline to which subsequent backup types are applied. For example, a full backup must be performed before a transaction log backup can be initiated.

Use the following steps to run a full backup:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
2. Right-click an **<Instance>** and click **Backup All Subclients**.
3. Select **Full** as the backup type and click **Immediate**.
4. Click **OK**.

You can track the progress of the job from the **Job Controller**. When the backup has completed, the **Job Controller** displays **Completed**.



TRANSACTION LOG BACKUPS

A transaction log backup captures the transaction log which contains a record of all committed or uncommitted transactions. Transaction log backups are consistent with the start time of the backup.

The use of transaction log backups make point-in-time recovery possible. This is useful in the scenario of a database failure where it is unacceptable to lose any data and you want to restore to the point of failure. If you use only full and differential backups, you will be able to restore to the time of the backup, but not to a point-in-time between backups.

A transaction log backup is similar to a traditional incremental backup you might perform on a file system because the transaction log backup contains only the new changes since the full or another transaction log backup.

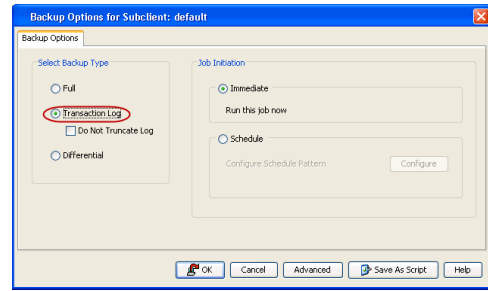
Each time a transaction log is backed up it is truncated to the exact time of the backup. No checkpoint is issued at this time, therefore dirty pages are not written to disk before or after a transaction log backup. If there are dirty pages, any completed transactions will need to be rolled forward if a transaction log restore is performed. Any transactions that are not completed at the time a transaction log backup is performed are rolled back during a restore involving a transaction log backup.

Use the following steps to run a transaction log backup:

1. Ensure that the SQL Server database is in full or bulk-logged recovery mode.
2. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server | <Instance>**.
3. Right click a **<Subclient>** and click **Backup**.
4. Select **Transaction Log** as backup type.

5. Click **Immediate**.
6. Click **OK**.

You can track the progress of the job from the **Job Controller**. When the backup has completed, the **Job Controller** displays **Completed**.

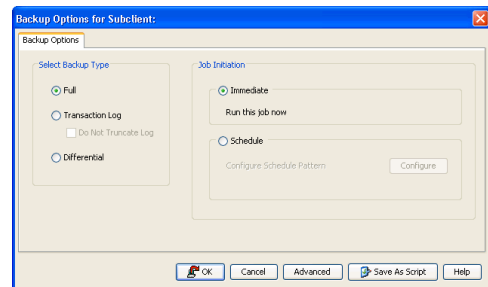


AUTOMATICALLY START A TRANSACTION LOG BACKUP AFTER A SUCCESSFUL BACKUP

You can start a Transaction Log backup automatically after a successful Full or Differential backup. This is useful when you want to back up logs immediately after a data backup, and allows you to do so without creating two scheduled jobs.

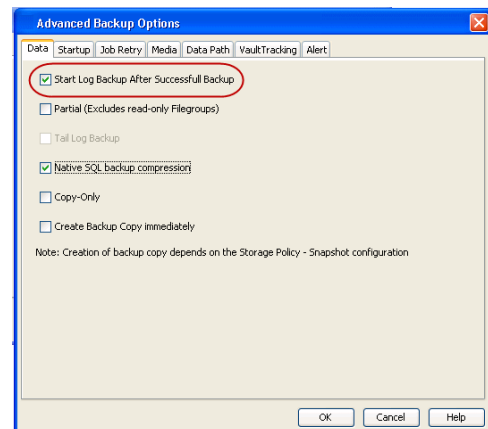
Use the following steps to automatically run a transaction log after a backup:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server | <Instance>**.
2. Right click an **<Subclient>** and click **Backup**.
3. Select **Full** or **Differential** backup type and click **Immediate**.



4. Click **Advanced**.
5. Click **Start Log Backup After Successful Backup**.
6. Click **OK**.

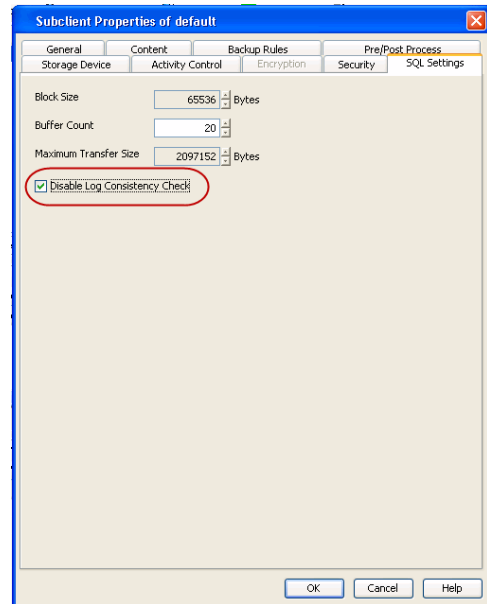
You can track the progress of the job from the **Job Controller**. When the backup has completed, the **Job Controller** displays **Completed**. Note that the simultaneous running of file or file group backup jobs is not supported.



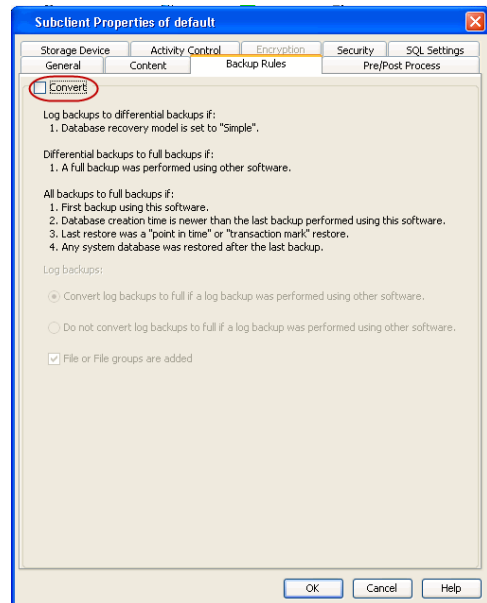
RUN TRANSACTION LOG BACKUPS WITHOUT RUNNING A FULL BACKUP

Use the following steps to perform transaction log backups without having to run full backups first.

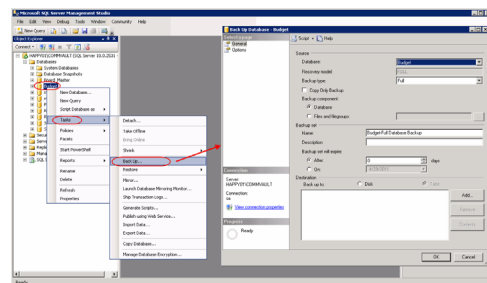
1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server | <Instance>**.
2. Right-click a **<Subclient>** and click **Properties**.
3. Click the **SQL Settings** tab and select **Disable Log Consistency Check**.
4. Click **OK**.



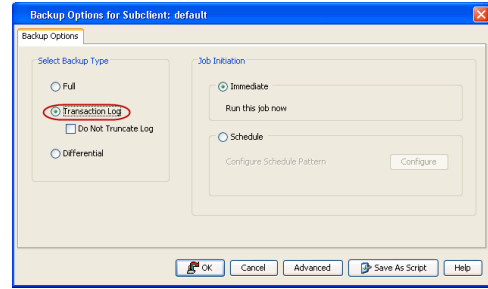
5. Click **Backup Rules** tab and clear **Convert** to disable the backup rules.
6. Click **OK**.



7. Run a full backup using SQL Server Management Studio (or any other application used to backup SQL Server).



8. Run a transaction log backup using the CommCell Console.



BACK UP TRANSACTION LOG OF DAMAGED DATABASE (DO NOT TRUNCATE LOGS)

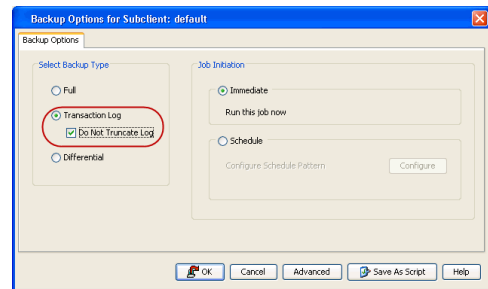
If you experience a database failure and you want to restore to the point of failure, a Transaction Log Backup with **Do not truncate log** must be initiated. This backups the database when it is damaged, regardless of its state.

It is used for capturing all transaction log events occurred since the last backup was run. This operation does not empty the active transaction log.

Use the following steps to disable log truncation during a backup:

1. Ensure that the SQL Server database is in full or bulk-logged recovery mode.
2. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server | <Instance>**.
3. Right click a <Subclient> and click **Backup**.
4. Select **Transaction Log** as backup type.
5. Select **Do Not Truncate Log**.
6. Click **Immediate**.
7. Click **OK**.

You can track the progress of the job from the **Job Controller**. When the backup has completed, the **Job Controller** displays **Completed**.



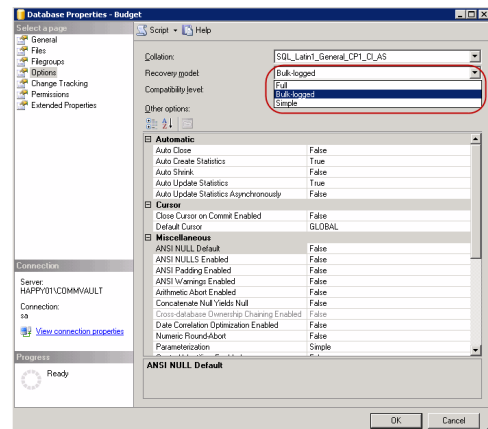
TAIL LOG BACKUP

When backing up transaction logs, you can choose to back up the tail of the log to capture the log records that have not yet been backed up. A tail-log backup prevents work loss and keeps the log chain intact. A tail-log backup allows you to recover a database to the point of failure; otherwise you can only recover a database to the end of the last backup that was created before the failure. For example, if a database was damaged or a data file was deleted, you should run a tail-log backup before attempting a file/file group restore. After the log tail is backed up, the database will be left in the RESTORING state.

Use the following step to backup the tail of a transaction log:

Ensure that the SQL Server database is in full or bulk-logged recovery model. To view or change the recovery model of a database:

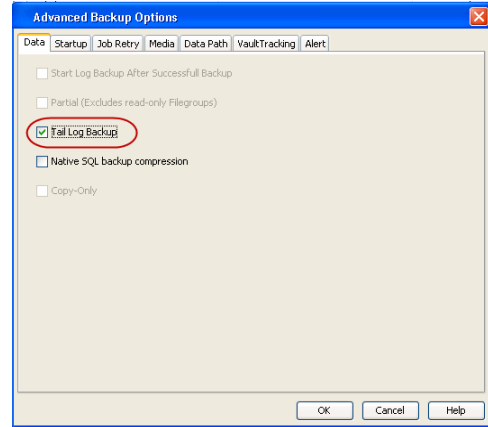
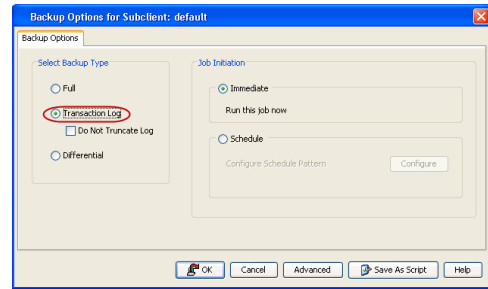
1. After connecting to the appropriate instance of the Microsoft SQL Server Database Engine, in Object Explorer, click the server name to expand the server tree.
2. Expand **Databases**, and, depending on the database, either select a user database or expand **System Databases** and select a system database.
3. Right-click the database, and then click **Properties**, which opens the **Database Properties** dialog box.
4. In the **Select a Page** pane, click **Options**.
5. The current recovery model is displayed in the **Recovery model** list box.
6. Select either **Full** or **Bulk-logged**.



7. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server | <Instance>**.
8. Right click a <Subclient> and click **Backup**.
9. Select **Transaction Log** as backup type.
10. Click **Immediate**.

11. Click **Advanced**.
12. Click **Tail Log Backup**.
13. Click **OK**.

You can track the progress of the job from the **Job Controller**. When the backup has completed, the **Job Controller** displays **Completed**. Note that the simultaneous running of file or file group backup jobs is not supported.



CONFIGURE NUMBER OF LOG BACKUPS BEFORE RUNNING A FULL BACKUP

Full backups are necessary at regular intervals as it reduces the chance of data loss if one of log backup becomes corrupted as it will invalidate (not restorable) all other log backups performed after that. This key is used for the purpose of re-enforcing the need of a full backup after certain number of transaction log backups have run.

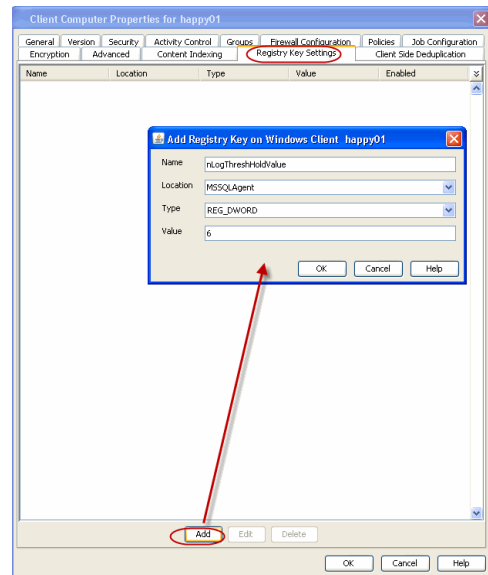
When this registry key is configured, a minor event will be generated in the Event Viewer to remind users to run a full backup after the configured number of transaction log backups have run.

Use the following steps to configure the number of log backups:

1. From the CommCell Browser, navigate to **Client Computers**.
2. Right-click the **<Client>** in which you want to add the registry key, and then click **Properties**.
3. Click the **Registry Key Settings** tab.
4. Click **Add**.
5. Enter `nLogThresholdValue` in the **Name** field.
6. Enter `MSSQLAgent` in the **Location** field.
7. From the **Type** list, select **REG_DWORD**.
8. Enter a number in the **Value** field. Range is [1 - <max_integer>].

This value specifies the number of transaction log backups that will be taken before a minor event is issued to remind users to run a full backup.

9. Click **OK**.

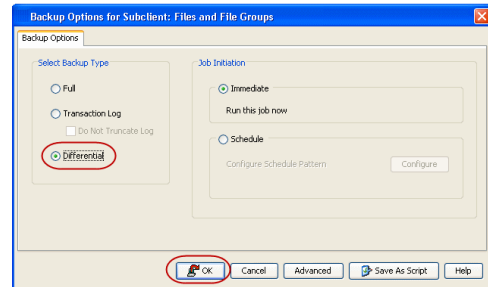


DIFFERENTIAL BACKUPS

A differential backup contains only the data that is new or has been changed since the last full backup. Differential backups consume less media and use less resources than full backups. Differential backups are cumulative. This means that each differential backup contains all changes accumulated since the last full backup. Each successive differential backup contains all the changes from the previous differential backup.

Use the following steps to run a differential backup:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server | <Instance>**.
2. Right-click a <Subclient> and click **Backup**.
3. Select **Differential** as the backup type and click **Immediate**.



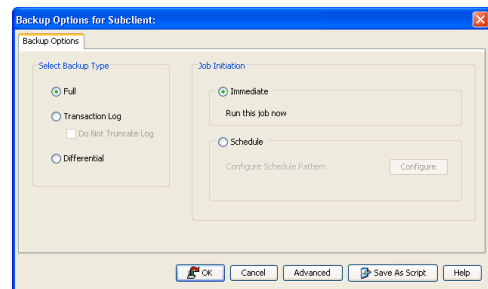
COMPRESSING BACKUPS WITH NATIVE SQL COMPRESSION (FOR SQL SERVER 2008 OR LATER)

Backups can be compressed before it is backed up to reduce the size of the backup. Typically, compressing a backup will require less device I/O which should increase backup speed significantly. However, CPU usage may increase for compressed backups and you may need to evaluate performance counters. Scheduling the backup during off-peak hours or compressing only low-priority backups may be desirable.

When using compression, there is no need for deduplication as the data will already be compressed and deduplication will not consequently save any more space.

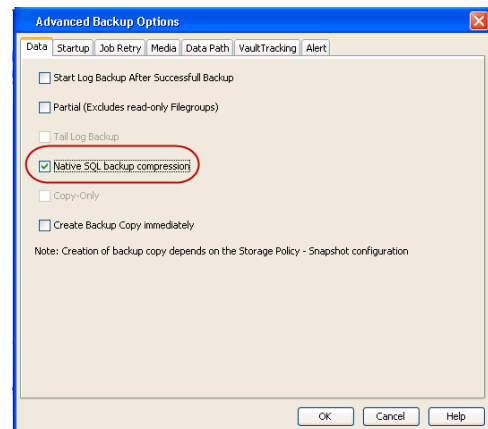
Use the following steps to enable compression:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server | <Instance>**.
2. Right-click a <Subclient> and click **Backup**.
3. Select a backup type and click **Immediate**.



4. Click **Advanced**.
5. Click **Native SQL backup compression**.
6. Click **OK**.

You can track the progress of the job from the **Job Controller**. When the backup has completed, the **Job Controller** displays **Completed**. Note that the simultaneous running of file or file group backup jobs is not supported.



PERFORM A PARTIAL BACKUP TO EXCLUDE READ-ONLY FILEGROUPS

A partial backup contains the following:

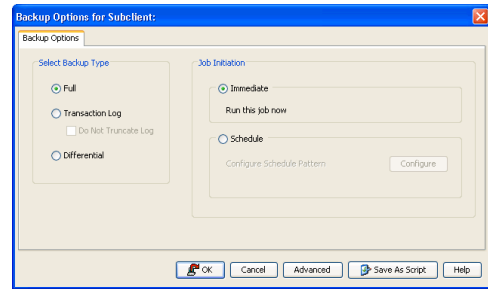
- All the data in the primary file group,
- Every read/write filegroup,
- Any optionally-specified read-only files

Partial backups are useful whenever you want to exclude read-only file groups. A partial backup is not supported when backing up transaction logs.

Use the following steps to enable partial backups:

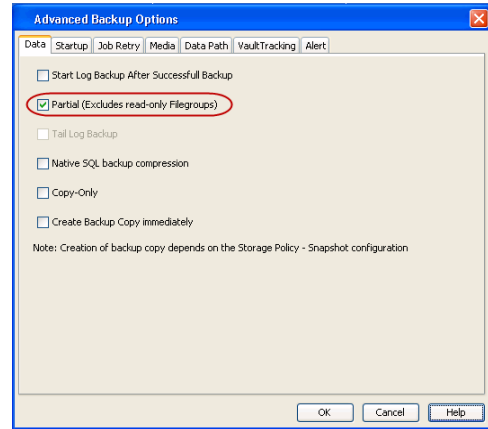
1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server | <Instance>**.

2. Right-click a <**Subclient**> and click **Backup**.
3. Select a backup type and click **Immediate**.



4. Click **Advanced**.
5. Click **Partial (Excludes read-only Filegroups)**.
6. Click **OK**.

You can track the progress of the job from the **Job Controller**. When the backup has completed, the **Job Controller** displays **Completed**. Note that the simultaneous running of file or file group backup jobs is not supported.



COMMAND LINE BACKUPS

You can perform backups of one or more SQL databases from the command line interface.

Command line backups enable you to perform backup operation on multiple clients simultaneously. In order to run the backups from command line, you need an input xml file which contains the parameters for configuring the backup options. This input xml file can be obtained from one of the following ways:

- Download the input xml file template and save it on the computer from where the backup will be performed.
- Generate the input xml file from the CommCell Console and save it on the computer from where the backup will be performed.

LOG ON TO THE COMMSERVE

To run command line operations you must first login to the CommServe as follows:

- From Command prompt, navigate to <Software_Installation_Directory>/Base and run the following command:

```
qlogin -cs <commserve name> -u <user name>
```

- For example, to log on to CommServe 'server1' with username 'user1':

```
qlogin -cs server1 -u user1
```

PERFORM THE BACKUP

1. Download the backup_template.xml file and save it on the computer from where the command will be executed.
2. Execute the following command from the <Software_Installation_Directory>/Base folder after substituting the parameters and attributes.

```
qoperation execute -af backup_template.xml -clientName xxxxx -instanceName xxxxx -subclientName xxxxx
```

3. Verify the status of the job using the following command:

```
qlist job -j JOBID
```

4. Once the job completes, logout from the CommServe using the qlogout command.

```
qlogout [-cs commserver] [-all] [-tf tokenfile] [-tk token] [-h]
```

EXAMPLES

Performing a Full Backup	<code>qoperation execute -af backup_template.xml -backupLevel FULL -subclientName subclient1 -clientName client1 -instanceName client1/instance1</code>
Performing a Transaction Log Backup	<code>qoperation execute -af backup_template.xml -backupLevel INCREMENTAL -subclientName</code>

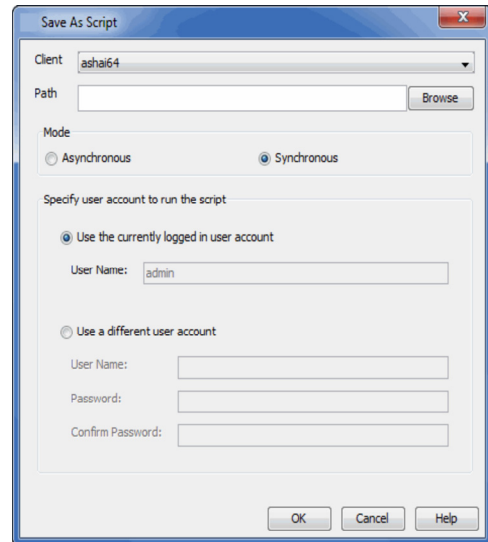

	subclient1 -clientName client1 -instanceName client1/instance1
Performing a Differential Backup	qoperation execute -af backup_template.xml -backupLevel DIFFERENTIAL -subclientName subclient1 -clientName client1 -instanceName client1/instance1
Performing an On Demand Backup	qoperation execute -af backup_template.xml -backupLevel FULL -subclientName subclient1 -clientName client1 -instanceName client1/instance1 -ondemandinputfile C:\test\myDBsContent.txt where myDBsContent.txt is an input file that list the databases as follows: DB1 DB2 To run ondemand backup for File File Group, the input file should list the database name, file group name and file name as follows: DB1<tab>Group1<tab>File1inGroup1 DB1<tab>Group2<tab>File2inGroup1

GENERATE THE COMMAND LINE SCRIPT FROM THE COMMCELL CONSOLE

In addition to the parameters provided in the template xml file, if you want to include additional options for the backup, you can do so by selecting the required options from the CommCell Console and generate the command line xml script for the backup.

Follow the steps given below to generate a script which you can use to perform a backup from the command line interface:


- From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server | <Instance>**.
- Right click an **<Subclient>** and click **Backup**.
- Select the required backup options which you want to execute using the script.
- Click **Save as Script**.
- Enter the location where you want to save the script or click **Browse** and navigate to the location.
 The script will be saved as a .xml file and a .bat file is created.
 If a file with the same name already exists in the specified location, the .xml file will be created with a timestamp. However, the .bat file will overwrite the existing file.
- Enter the username and password for the user account which you want to use to perform the backup.
 By default, the user account which you have used to login to CommCell console is used for performing the backup. However, if the user account does not have access to any application or database, click **Use a different account**.
- Click **OK**.



SCHEDULING A BACKUP

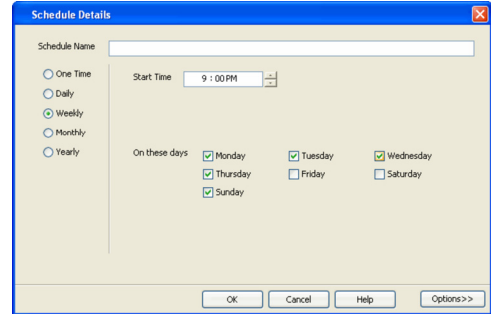
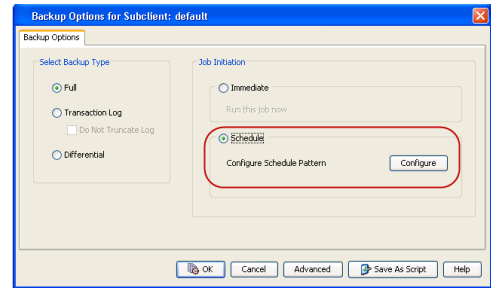
Follow the steps given below to schedule a backup:

- From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server | <Instance>**.
 - Right-click a **<Subclient>** and click **Backup**.
- Select the **Backup type**.
 - Click **Schedule** to schedule the backup for a specific time.
 - Click **Configure** to set the schedule for the backup job. The Schedule Details dialog displays.



3. Select the appropriate scheduling options. For example:
 - Click **Weekly**.
 - Check the days you want the run the backup job.
 - Change the Start Time to 9:00 PM.
 - Click **OK** to close the Schedule Details dialog.
 - Click **OK** to close the Backup Options dialog.

The backup job will execute as per the schedule.



See Scheduling for a comprehensive information on scheduling jobs.

MANAGING JOBS

Jobs can be managed in a number of ways. The following sections provide information on the different job management options available:

RESTARTING JOBS

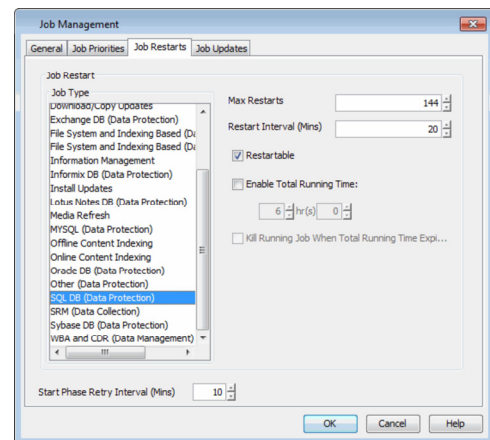
Jobs that fail to complete successfully are automatically restarted based on the job restartability configuration set in the Control Panel. Keep in mind that changes made to this configuration will affect all jobs in the entire CommCell.

To Configure the job restartability for a specific job, you can modify the retry settings for the job. This will override the setting in the Control Panel. It is also possible to override the default CommServe configuration for individual jobs by configuring retry settings when initiating the job. This configuration, however, will apply only to the specific job.

Backup jobs for this Agent are resumed from the point-of-failure.

CONFIGURE JOB RESTARTABILITY AT THE COMMSERVE LEVEL

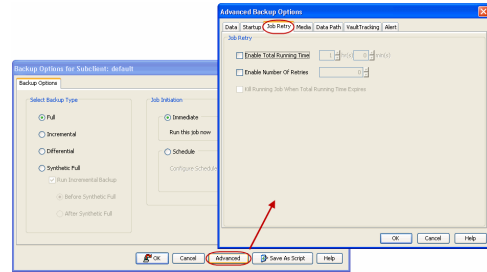
1. From the CommCell Browser, click **Control Panel** icon.
2. Select **Job Management**.
3. Click **Job Restarts** tab and select a **Job Type**.
 - o Select **Restartable** to make the job restartable.
 - o Change the value for **Max Restarts** to change the maximum number of times the Job Manager will try to restart a job.
 - o Change the value for **Restart Interval (Mins)** to change the time interval between attempts for the Job Manager to restart the job.
4. Click **OK**.



CONFIGURE JOB RESTARTABILITY FOR AN INDIVIDUAL JOB

1. From the CommCell Console, navigate to **<Client> | SQL Server | <Instance>**.
2. Right-click the **Subclient** and select **Backup**
3. Click **Advanced**.

4. In the **Advanced Backup Options** dialog box, click the **Job Retry** tab.
5. Select **Enable Total Running Time** and specify the maximum elapsed time before a job can be restarted or killed.
 Select **Kill Running Jobs When Total Running Time Expires** to kill the job after reaching the maximum elapsed time.
6. Select **Enable Number Of Retries** and specify the number of retries.
7. Click **OK**.



CONTROLLING JOBS

The following controls are available for running jobs in the Job Controller window:

SUSPEND	Temporarily stops a job. A suspended job is not terminated; it can be restarted at a later time.
RESUME	Resumes a job and returns the status to Waiting, Pending, Queued, or Running. The status depends on the availability of resources, the state of the Operation Windows, or the Activity Control setting.
KILL	Terminates a job.

SUSPENDING A JOB

1. From the Job Controller of the CommCell Console, right-click the job and select **Suspend**.
2. The job status may change to **Suspend Pending** for a few moments while the operation completes. The job status then changes to **Suspended**.

RESUMING A JOB

1. From the Job Controller of the CommCell Console, right-click the job and select **Resume**.
2. As the Job Manager attempts to restart the job, the job status changes to **Waiting, Pending, or Running**.

KILLING A JOB

1. From the Job Controller of the CommCell Console, right-click the job and select **Kill**.
2. Click **Yes** when the confirmation prompt appears if you are sure you want to kill the job. The job status may change to **Kill Pending** for a few moments while the operation completes. Once completed, the job status will change to **Killed** and it will be removed from the Job Controller window after five minutes.

See Job Management for a comprehensive information on managing jobs.

ADDITIONAL OPTIONS

The following table describes the available additional options to further refine your backup operations:

OPTION	DESCRIPTION	RELATED TOPICS
Startup Options	<p>The Startup Options are used by the Job Manager to set priority for resource allocation. This is useful to give higher priority to certain jobs. You can set the priority as follows:</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <Client> SQL Server <Instance>. 2. Right-click the Subclient in the right pane and click Backup. 3. Click Advanced and click Startup tab. 4. Select the Change Priority checkbox. 5. Enter a priority number - 0 is the highest priority and 999 is the lowest priority. 6. Select the Start up in suspended State check box to start the job in a suspended state. 7. Click OK. 	Refer to Job Priority and Priority Precedence.
Alerts	<p>This option enables users or user groups to get automatic notification on the status of the data protection job. Follow the steps given below to set up the criteria to raise notifications/alerts:</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <Client> SQL Server <Instance>. 2. Right-click the Subclient in the right pane and click Backup. 3. Click Advanced and select the Alert tab. 4. Click Add Alert. 	Refer to Alerts.

	<ol style="list-style-type: none"> 5. From the Add Alert Wizard dialog box, select the required threshold and notification criteria and click Next. 6. Select the required notification types and click Next. 7. Select the recipients and click Next. 8. Click Finish. 9. Click OK. 	
Command Line Backups	<p>Command Line Interface enables you to perform backups or restore from the command line. The commands can be executed from the command line or can be integrated into scripts.</p> <p>You can also generate command line scripts for specific operations from the CommCell Browser using the Save As Script option.</p>	Refer to Command Line Interface.
CommCell Readiness Report	<p>The CommCell Readiness Report provides you with vital information, such as connectivity and readiness of the Client, MediaAgent and CommServe. It is useful to run this report before performing the data protection or recovery job. Follow the steps given below to generate the report:</p> <ol style="list-style-type: none"> 1. From the Tools menu in the CommCell Console, click Reports. 2. Navigate to Reports CommServe CommCell Readiness. 3. Click the Client tab and click the Modify button. 4. In the Select Computers dialog box, clear the Include All Client Computers and All Client Groups check box. 5. Select the client from the Exclude list. 6. Click the Include > button. 7. Click OK. 8. Click the MediaAgent tab. 9. Clear the Include All MediaAgents checkbox. 10. Select the MediaAgent from the Exclude list. 11. Click Include >. 12. Click Run. <p>The generated report is displayed.</p>	Refer to CommCell Readiness Report.
Backup Job Summary Report	<p>The Backup Job Summary Report provides you with information about all the backup jobs that are run in last 24 hrs for a specific subclient. You can get information such as status, time, data size etc. for each backup job. It is useful to run this report after performing the backup. Follow the steps given below to generate the report:</p> <ol style="list-style-type: none"> 1. From the Tools menu in the CommCell Console, click Reports. 2. Navigate to Reports Jobs Job Summary. 3. Click Data Management on the General tab in the right pane. 4. Select the Computers tab. 5. Click Subclient and select the Edit tab. 6. Navigate to Client Computers <Client> File System Backup Set Subclient. 7. Click Run. 	Refer to Backup Job Summary Report.
Data Path Options	<p>Data Protection operations use a default Library, MediaAgent, Drive Pool, and Drive as the Data Path. You can use this option to change the data path if the default data path is not available. Follow the steps given below to change the default data path:</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <Client> SQL Server <Instance>. 2. Right-click the Subclient in the right pane and click Backup. 3. Click Advanced and select the Data Path tab. 4. Select the MediaAgent and Library. 5. Select the Drive Pool and Drive for optical and tape libraries. 6. Click OK. 	Refer Change Data Path.
Start New Media	<p>The Start New Media option enables you to start the data protection operation on a new media. This feature provides control over where the data physically resides. Use the following steps to start the data protection operation on a new media:</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <Client> SQL Server <Instance>. 2. Right-click the Subclient in the right pane and click Backup. 	Refer to Start New Media.

	<ol style="list-style-type: none"> 3. Click Advanced and select the Media tab. 4. Select the Start New Media check box. 5. Click OK. 	
Mark Media Full on Success	<p>This option marks the media as full, two minutes after the successful completion of the data protection job. This option prevents another job from writing to this media. Follow the steps given below:</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <Client> SQL Server <Instance>. 2. Right-click the Subclient in the right pane and click Backup. 3. Click Advanced and select Media tab. 4. Select the Mark Media Full on Success check box. 5. Click OK. 	Refer to Export Media.
Allow other Schedules to use Media Set	<p>The Allow Other Schedules to use Media Set option allows jobs that are part of the schedule or schedule policy and using the specific storage policy to start a new media. It also prevents other jobs from writing to the same set of media.</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <Client> SQL Server <Instance>. 2. Right-click the Subclient in the right pane and click Backup. 3. Click Advanced and select the Media tab. 4. Select the Allow Other Schedules To Use Media Set check box. 5. Click OK. 	Refer to Creating an Exportable Media Set.
Extended Data Retention	<p>This option allows you to extend the expiration date of a specific job. This will override the default retention set at the corresponding storage policy copy. Follow the steps given below to extend the expiration date:</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <Client> SQL Server <Instance>. 2. Right-click the Subclient in the right pane and click Backup. 3. Click Advanced and select the Media tab. 4. Select one of the following options: <ul style="list-style-type: none"> o Infinite - Select this option to extend the expiration date by infinite number of days o Number of day - Select this option to specify the number of days to extend the expiration date and then enter the number of days. 5. Click OK. 	Refer to Extended Retention Rules.
Vault Tracker	<p>This feature provides the facility to manage media that is removed from a library and stored in offsite locations. Depending on your VaultTracker setup, select the required options. Use the following steps to access and select the VaultTracker options.</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <Client> SQL Server <Instance>. 2. Right-click the Subclient in the right pane and click Backup. 3. Click Advanced and select the VaultTracking tab. 4. Select the required options. 5. Click OK. 	Refer to VaultTracker or VaultTracker Enterprise.

Browse Data – SQL Server iDataAgent

TABLE OF CONTENTS

Understanding the Browse Window

Browsing Data

- Latest Data
- Data Before a Specified Time
- Data Between a Specified Time


Browse from Copies

List Media

UNDERSTANDING THE BROWSE WINDOW

Browse provides an easy-to-use convenience to search and restore data.

You can open the browse window from client, agent, instance or subclient level. The sub levels displayed in the browse window is based on where you start the browse operation.



SELECTING OBJECTS FROM THE BROWSE WINDOW FOR RESTORE

The browse window displays objects and consists of two parts:

- The left pane displays the object tree at the selected level.
- The right pane displays the contents of the selected object.

Note that the window displays only the data that was obtained by a backup. Data that is excluded by a filter, or data which did not exist as of the specified browse time does not appear in the window.

Selections follow these rules:

- All selections are recursive.
- Clicking a box again causes the selection to be cleared.
- If you select an object in the left pane, then all of its contents are selected.
- You can select a specific object in the right pane.

Selection status is revealed by the selection icons as follows:

<input type="checkbox"/>	Indicates that the object is not selected for restoration.
<input checked="" type="checkbox"/>	Indicates that a portion of the object is selected for restoration. i.e., only some of the child object(s) within the selected object.
<input checked="" type="checkbox"/>	Indicates that the entire object, including all of its child objects, are selected for restoration.

BROWSING DATA

The option to browse the backup data provides the facility to view and restore the data that was backed up. The following sections explain how to browse the backup data.


LATEST DATA

By default, you can browse the latest backup data.

Follow the steps given below to view the latest data backed up by a backup set:

- From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
- Right-click the **<Instance>** and then click **All Tasks | Browse Backup Data**.
- Click **OK**.

The latest data backed up by all the subclients is displayed in the **Client Browse** tab.



DATA BEFORE A SPECIFIED TIME

Follow the steps given below to browse the data backed up before specified time:

- From the CommCell Browser, navigate to **Client Computers | <Client> | SQL**

RELATED TOPICS

Scheduling


Provides comprehensive information on scheduling jobs.

Job Management

Provides comprehensive information on managing jobs.

Server.

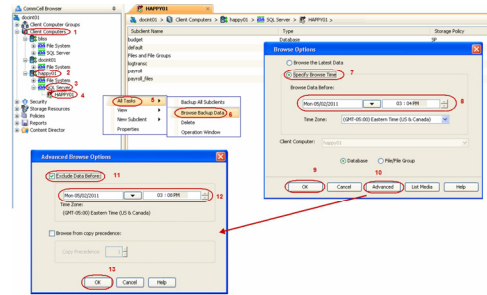
2. Right-click the instance and then click **All Tasks | Browse Backup Data**.
3. Select **Specify Browse Time**.
4. Select a date and time to **Browse Data Before** it.
5. Click **OK**.



DATA BETWEEN A SPECIFIED TIME

Follow the steps given below to browse data between specified time:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | File System**
2. Right-click a Backup Set and click **All Tasks | Browse/Browse Backup Data**.
3. Select **Specify Browse Time**.
4. Select a date and time to **Browse Data Before** it.
5. Click **Advanced**.
6. In the **Advanced Browse Options** dialog box, select **Exclude Data Before**.
7. Select the date and time to exclude the data before it from the browse operation.
8. Click **OK** in the **Advanced Browse Options** dialog box .
9. Click **OK** in the **Browse Options** dialog box.




BROWSE FROM COPIES

By default, when a browse operation is requested, the software attempts to browse from the storage policy copy with the lowest copy precedence. If the data that you want to browse was already pruned from the primary copy, the software searches the other copies of the storage policy, starting from the copy with the lowest copy precedence to a copy with the highest copy precedence.

This feature is useful in the following conditions:

- The media containing protected data for a particular copy has been removed from the storage library, you can choose to browse from a copy whose media are inside the library.
- Allows browsing from a copy that accesses faster disk media rather than slower tape media.
- When media drives used by a particular copy are busy with another operation, this helps in browsing from a different copy to avoid resource conflicts.

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
2. Right-click the instance and then click **All Tasks | Browse Backup Data**.
3. From the **Browse Options** dialog box, click **Advanced**.
4. In the **Advanced Browse Options** dialog box select the **Browse from copy precedence** option.



If you specify a copy precedence number for a browse operation, the software searches only the storage policy copy with that precedence number in all storage policies used for securing the data. If data does not exist in the specified copy, the browse operation fails even if the data exists in another copy of the same storage policy.

5. Specify the precedence number in **Copy Precedence**.
6. Click **OK** in the **Advanced Browse Options** dialog box .
7. Click **OK** in the **Browse Options** dialog box.

LIST MEDIA

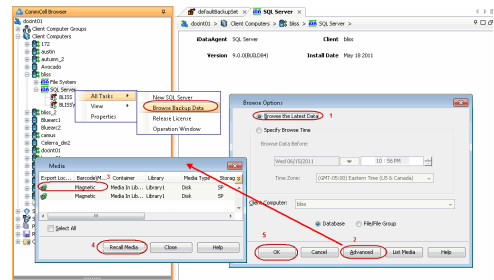
List media option is useful to predict media required to restore the index required to browse data. This is useful when the index is not available in the index cache.

The following section describes how to perform this operation.

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.

2. Right-click the instance and then click **All Tasks | Browse Backup Data**.
3. From the **Browse Options** dialog box, if required, select the following options:
 - o Click **Browse the Latest Data** to list media associated with the most recent data protection cycle.
 - o Click **Specify Browse Time** to list media associated with data protection operations up to the specified date and time range. Use the **Browse Data Before** box to specify the end date and time.
 - o Click **Advanced** and then click **Exclude Data Before** and then select the date and time from which you wish to list media associated with data protection operations.
4. Click **List Media**.
5. From the **List Media** dialog box, select the media you wish to recall and click **Recall Media**.
6. From the **Recall Media** dialog box General tab, select the following:
 - o The time until which the media would be retained in the library for read operations.
 - o A reason for recalling the media.
7. From the **Destination** tab, configure the following destination options:
 - o Click the **Track Transit** option and select the transit location from the list, to track the transit information.
 - o Select the desired **Destination**.
 - o If desired, select the **Move Media to Overwrite Protection Pool** option along with the desired pool to which the media will be moved.
 - o Select **Acknowledge the action as Reached Destination automatically** if desired.
8. Click **OK**.

The appropriate media is listed in the **Media** dialog box.



Replication Using Warm Database Restore

TABLE OF CONTENTS

Replication Using Warm Database Restore

- Pre-Requisites
- Setting Up the Standby Server
- Using the SQL Server Database(s) on the Hot Standby Server
- Performing Disaster Recovery Operations for the Primary SQL Server Database(s)

REPLICATION USING WARM DATABASE RESTORE

This type of restore is also called Hot Server Restore. You can keep up-to-date copies of the production databases on alternate servers, thus minimizing down time in cases where the production database server has failed.

The advantages of Warm Database Restore over a traditional restore scenario is that the database is always kept in a near ready state. In the case of disaster, users do not need to wait for the time it would take to restore the database in its entirety. Only the latest backup may need to be applied to the target database before turning the application over to use it.

PRE-REQUISITES

Review the following requirements before performing a Warm Database Restore:

- The primary SQL server and hot standby server must be clients of the same CommServe and have the SQL Server iDataAgent installed on each computer.
- The primary SQL server must be able to communicate with the hot standby server with one of the following network configurations:
 - Local Area Network (LAN) in the same domain
 - Local Area Network (LAN) in a different domain
 - Wide Area Network (WAN)
- Two licenses for SQL Server iDataAgent is required.

SETTING UP THE STANDBY SERVER

Use the following steps to setup the standby server:

1. Install the Microsoft SQL Server iDataAgent on the hot standby server computer.
2. Perform a full backup of the SQL Server databases on the primary server.
3. Perform a restore of the last full backup of the primary database, as well as any differential or transaction log backups completed since the database's last full backup.
4. Schedule routine differential or transaction log backups of the SQL Server databases on the primary server.
5. Schedule routine restore operations of the **Latest Backup Data** in the Standby Server.
 - Matching the restore schedule with a backup schedule on the source database automatically keeps the restored databases in the Standby Server up-to-date.
 - If you wish to include only the transaction log backups completed since the last restore operation, you must select the **Apply Log Backups Only** option after selecting **Latest Backup Data** in the **SQL Restore Options** dialog box. The restore operation will then skip any full backup jobs completed since the last restore operation.

If you do not select the **Apply Log Backups Only** option and the database already exists, all backup jobs completed since the last restore operation will automatically be restored.

If you do not select the **Apply Log Backups Only** option and the database does not already exist, the software will automatically perform a full restore of the database.

USING THE SQL SERVER DATABASE(S) ON THE HOT STANDBY SERVER

In the event of a failure in the primary server you can use the hot standby server by bringing the databases online. Use the following steps to bring the database online.

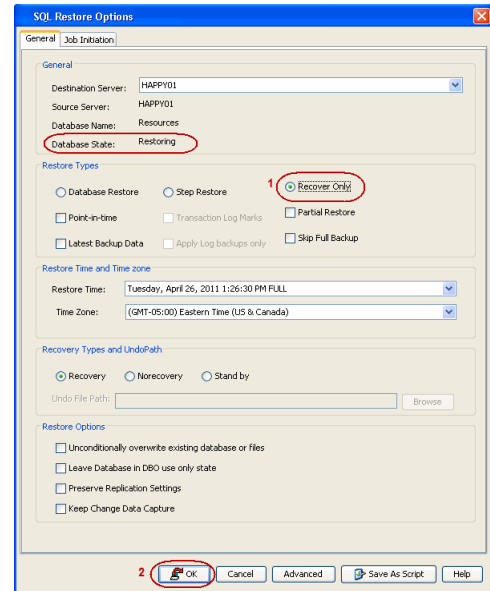
1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
2. Right-click the instance and then click **All Tasks | Browse Backup Data**.
3. Click **OK**.
4. In the **Browse** window, select the databases that you want to recover in the right pane and click **Recover All Selected**.

5. Click **Recover Only**.
6. Click **OK**.

You can also bring the databases online by recovering the database from the **SQL Server Enterprise Manager**.

No further restore operations may be performed to the hot standby server's databases once the databases have been recovered.

If you wish to perform more restore operations to the hot standby server after recovering the databases, you must setup the hot standby server once again.



PERFORMING DISASTER RECOVERY OPERATIONS FOR THE PRIMARY SQL SERVER DATABASE(S)

If the primary server's databases is lost, you can restore the databases from the hot standby server's databases by performing the appropriate in-place restore operation.

If the entire primary server is lost, you must perform a full system restore.

Once the full system restore has completed, you will need to set up the hot standby server again by repeating the steps outlined in this procedure.

Advanced Restore - SQL Server iDataAgent

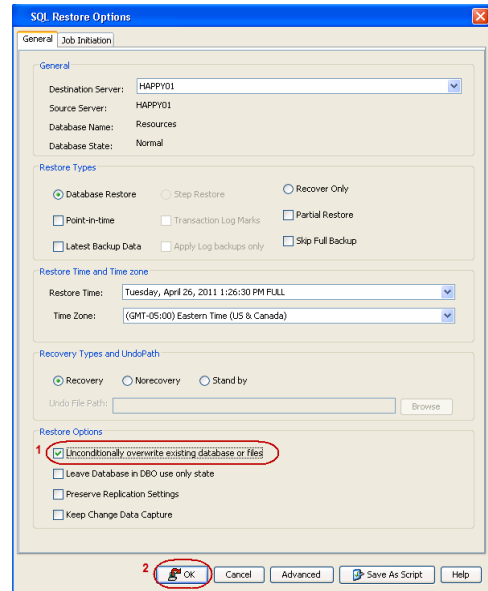
TABLE OF CONTENTS

<p>Restoring a Database</p> <p>Recovering a Database</p> <ul style="list-style-type: none"> Recover a Database in the Offline State Recover a Database in the Standby State Recover one or more Databases in the Online State <p>Restoring System Databases</p> <p>Restoring Files or Filegroups</p> <ul style="list-style-type: none"> Restore One or More Files or Filegroups Restore a Database using File/Filegroup Level Performing Partial Restore of a Database <p>Restoring SQL Server Transaction Logs (Step Restore)</p> <p>Applying Log Backups</p> <p>Restoring SQL Databases to a Point in Time</p> <p>Restoring Only the Logs</p> <p>Restoring SQL Server Databases to a Transaction Mark</p> <p>Restoring a Database With Restricted Access</p> <p>Preserving Replication Settings</p> <p>Keeping Change Data Capture (For SQL Server 2008 and Above)</p> <p>Increasing the VDI Timeout</p> <p>Performing an Out-of-Place Restore</p> <ul style="list-style-type: none"> Copying a Database to a Different Location Copying a Database to a Different SQL Server Instance Moving a Database Moving a Database Back to the Original Location After a Move <p>Scheduling a Restore</p> <p>Setting Up Pre-Post Processes</p> <ul style="list-style-type: none"> Setting Up a Pre/Post Process to Run After Each Attempt <p>Managing Restore Jobs</p> <ul style="list-style-type: none"> Restarting Jobs Controlling Jobs <p>Additional Restore Options</p>	<p>Command Line Restores</p> <ul style="list-style-type: none"> Log on to the CommServe Perform the Restore Generate the Command Line Script from the CommCell Console
---	--

RESTORING A DATABASE

By default, a database is restored in the same location from where it was backed up and the existing database files are overwritten. This restore leaves the database in an online state. Follow the steps given below to restore a database:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
2. Right-click the **<Instance>** and then click **All Tasks | Browse Backup Data**.
3. Click **OK**.
4. In the right pane of the Browse window, select a non-system database you want to restore and click **Recover All Selected**.
5. Click **Unconditionally overwrite existing database or files**.
6. Click **OK** to start the restore.



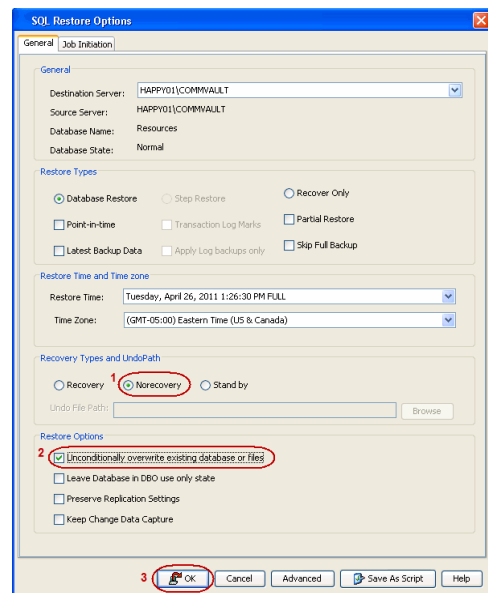
RECOVERING A DATABASE

By default, a database is left in the online state after a restore. However, you might need to change the state of the database depending on your needs (e.g., if you want to make the database inaccessible to the users). By recovering a database, you can select the state in which the database is to be left. Follow the steps given below to recover a database:

RECOVER A DATABASE IN THE OFFLINE STATE

If you want the database to be offline after the restore, follow the steps given below:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
2. Right-click the **<Instance>** and then click **All Tasks | Browse Backup Data**.
3. Click **OK**.
4. In the right pane of the Browse window, select the database to be recovered and click **Recover All Selected**.
5. Click **Norecovery**.
6. Click **Unconditionally overwrite existing database or files**.
7. Click **OK**.



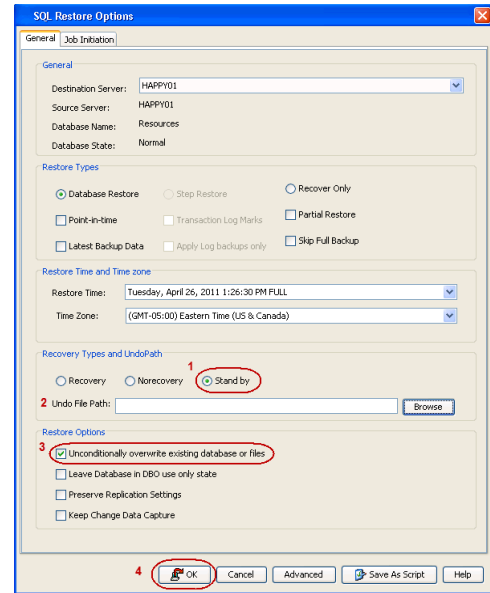
RECOVER A DATABASE IN THE STANDBY STATE

If you want the database to be in a read-only state after the restore, follow the steps given below:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
2. Right-click the **<Instance>** and then click **All Tasks | Browse Backup Data**.
3. Click **OK**.
4. In the right pane of the Browse window, select the database to be recovered and click

Recover All Selected.

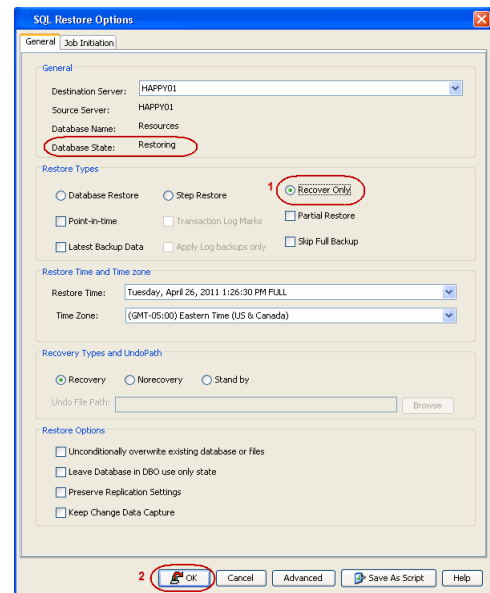
5. Click **Stand by**.
Specify the path to the undo file or click **Browse** to locate the file.
6. Click **Unconditionally overwrite existing database or files**.
7. Click **OK**.



RECOVER ONE OR MORE DATABASES IN THE ONLINE STATE

If you have a database in the offline or standby state and you want to bring the database back to the online state, follow the steps given below:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
2. Right-click the **<Instance>** and then click **All Tasks | Browse Backup Data**.
3. Click **OK**.
4. In the right pane of the Browse window, select the databases you want to recover and click **Recover All Selected**.
5. Click **Recover Only**.
6. Click **OK**.



RESTORING SYSTEM DATABASES

System databases are essential for the operation of the SQL server instance. If you need to restore the SQL Server in the event of a system failure, the system databases (*master*, *msdb* and *model*) must be backed up. The *tempdb* does not get backed up as it is re-created by the SQL Server every time the server is started.

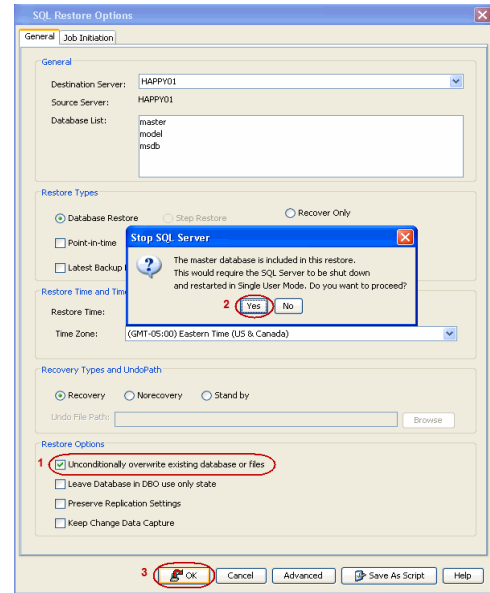
By default, system databases are restored in the same location from where they were backed up and their data files are overwritten. Follow the steps given below to restore the system databases:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
2. Right-click the instance, point to **All Tasks**, and then click **Browse Backup Data**.
3. Click **OK**.
4. In the left pane of the **Browse** window, navigate to the instance node.
5. Select the system databases (*master*, *msdb*, and *model*) in the right pane and click

Recover All Selected.

If you choose to restore the system databases individually, it is recommended to follow this order: master, msdb, model.

6. Click **Unconditionally overwrite existing database or files.**
7. Click **OK** to start the restore.
8. Click **Yes.** The system stops and restarts the SQL Server service in single-user mode.



RESTORING FILES OR FILEGROUPS

By default, files or filegroups are restored in the same location from where they were backed up. File/Filegroup restores give you the ability to bring back the files that have been damaged in the event of a failure. You can restore specific files and filegroups from both:

- database backups
- file and filegroup backups

During a File/Filegroup restore job, the system restores the latest full and differential backups and then uses the transaction log backups up to the specified restore time. Hence, transaction logs are needed to successfully restore files or filegroups as it determines the end time of the restore. For more information, see Transaction Log Backups.

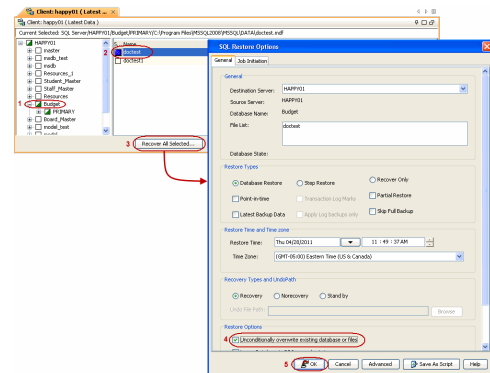
File/Filegroup restores are commonly used for situations in which individual data files on a database are damaged.

RESTORE ONE OR MORE FILES OR FILEGROUPS

For a single data file, the file/file group restore chain consists of the most recent full backup, the most recent differential (if any), and all subsequent transaction log backups that occurred prior to the restore time. However, for multiple files or file groups, the same rules apply for each data file, and the log backups to be restored will be determined by the file requiring the oldest log.

Follow the steps below to restore one or more files or filegroups from a database:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server.**
2. Right-click the **<Instance>** and then click **All Tasks | Browse Backup Data.**
3. Click **File/File Group** and then click **OK.**
4. In the left pane of the Browse window, navigate to the database that contains the files or filegroups you want to restore.
5. Select the files or filegroups you want to restore in the right pane and click **Recover All Selected.**
6. Click **Unconditionally overwrite existing database or files.**
7. Click **OK** to start the restore.



RESTORE A DATABASE USING FILE/FILEGROUP LEVEL

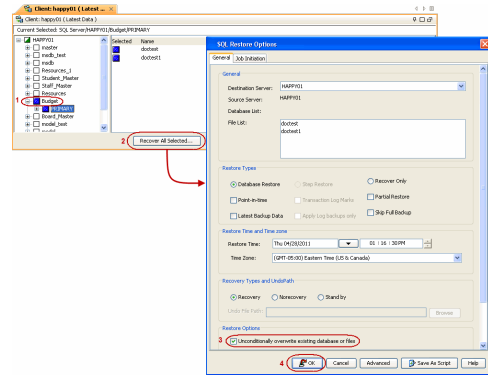
You can restore a database in its entirety by restoring all the filegroups within the database. Follow the steps below to restore the entire database at the File/filegroup level:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server.**

2. Right-click the **<Instance>** and then click **All Tasks | Browse Backup Data**.
3. Click **File/File Group** and then click **OK**.
4. In the right pane of the Browse window, select the database you want to restore.

Multiple databases cannot be restored at the filegroup level.

5. Click **Recover All Selected**.
6. Click **Unconditionally overwrite existing database or files**.
7. Click **OK** to start the restore.



PERFORMING PARTIAL (PIECEMEAL) RESTORE OF A DATABASE

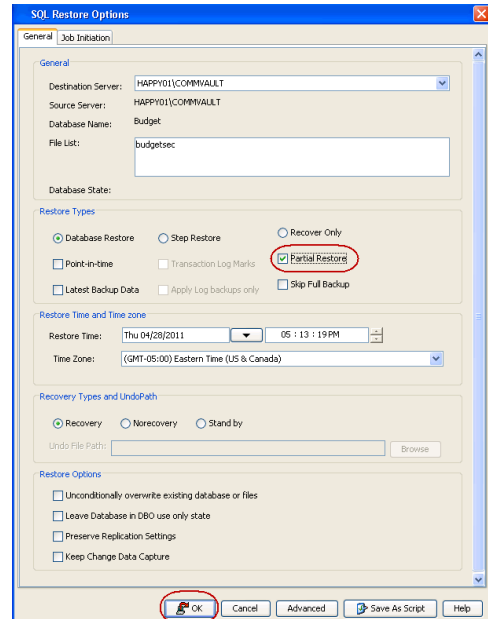
If the size of a filegroup in a database is large, the restore operation may take considerable time. In such case, you can restore the database in stages.

Partial restores also known as Piecemeal Restore in SQL Server versions 2005 and later allows you to restore a database in stages.

Follow the steps given below to restore a database in stages at filegroup level:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
2. Right-click the **<Instance>** and then click **All Tasks | Browse Backup Data**.
3. Click **File/File Group** and then click **OK**.
4. In the left pane of the Browse window, navigate to the database that contains the filegroups you want to restore.
5. Select the filegroups you want to restore in the right pane and click **Recover All Selected**.
6. Select the **Partial Restore** check box.
7. Click **OK** to start the restore.
8. Perform the partial restore of all remaining filegroups one by one to restore the entire database.

After performing a partial (piecemeal) restore job, the next backup job for that database is automatically converted to a full backup.



RESTORING SQL SERVER TRANSACTION LOGS (STEP RESTORE)

When you want to restore an event in the database but do not know the exact time of the event, you can first restore the database in the standby state and apply transaction logs to the database. The Step Restore enables you to append the logs to such a database. (Both Incremental or Transaction Log backups can be applied.) Use the following procedure to apply the log backups one by one to scan through the changes that have been committed between each log backup. This way you can avoid performing multiple full restores when you are not sure of the point of restore required.

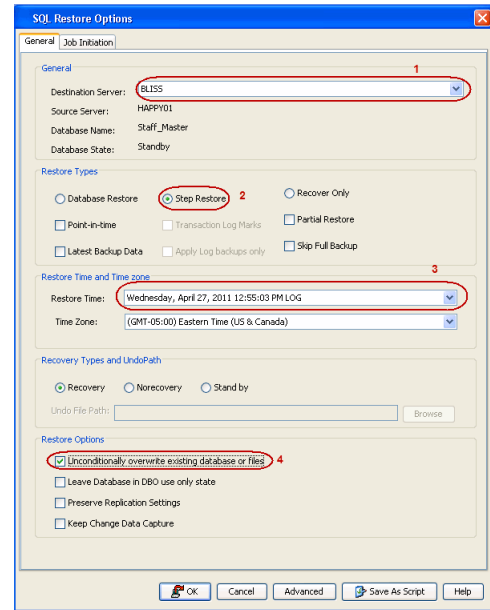
Step Restores can be performed on individual databases.

Follow the steps given below to apply the logs to a database in the standby state:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
2. Right-click the **<Instance>** and then click **All Tasks | Browse Backup Data**.
3. Click **OK**.
11. In the right pane of the Browse window, select the database that is already restored in the standby state and click **Recover All Selected**.
12. Select the **Destination Server**. The destination Server must be different than source server.
13. Select **Step Restore**.

14. Select a **Log** backup from the **Restore Time** list.
15. Select the **Unconditionally overwrite existing database or files** check box.
16. Click **OK** to start the log restore.

If required, you can apply logs from other Log backups by selecting a different Log backup from the Restore Time list.

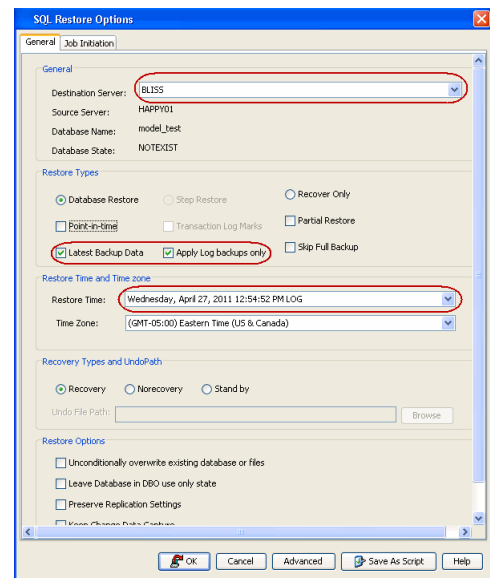


APPLYING LOG BACKUPS

This option enables you to apply all the latest transaction logs to a hot standby sever. The hot standby server is the server where you can restore a latest backup data periodically. Follow the steps given below to restore only the logs that are backed on a specific time:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
2. Right-click the **<Instance>** and then click **All Tasks | Browse Backup Data**.
3. Click **OK**.
4. In the right pane of the Browse window, select a database that is already restored in the standby state and click **Recover All Selected**.
5. Select the **Destination Server**. Select the hot standby server.
6. Select the **Latest Backup Data** check box.
7. Select the **Apply Log backups only** check box.
8. Select the time and a **Log** backup from the **Restore Time** list.
9. Click **OK** to start the log restore.

The Transaction Logs, backed up on the selected time will be restored and applied to the database.



RESTORING SQL DATABASES TO A POINT IN TIME

If any undesired transaction occurs in the database, you can revert the database to a state just before the transaction. The point in time option enables you to restore the database to a specific point-in-time. This option is useful in the following scenarios:

- If any undesired transaction occurs in the database, you can revert the database to a state just before the transaction.
- If a database fails, you can restore to the state just before the point of failure.
- You can restore multiple databases to a consistent time. This will be useful for the absolute synchronization of the databases.

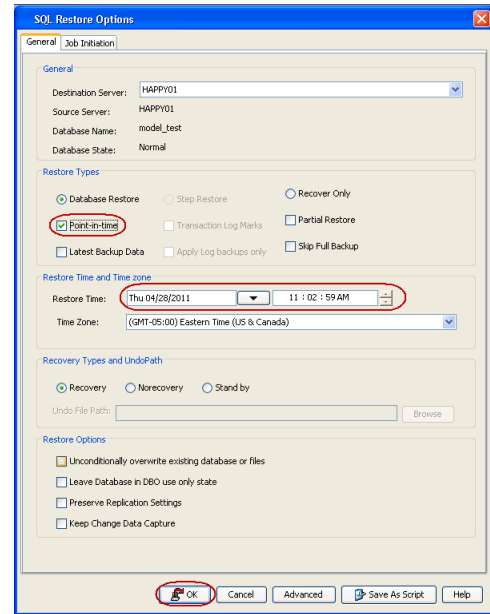
Before performing the point-in-time restore of a database, ensure that transaction logs for the database are backed up. Follow the steps given below to restore a database(s) to a point in time:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
2. Right-click the **<Instance>** and then click **All Tasks | Browse Backup Data**.

3. Click **OK**.
4. In the right pane of the Browse window, select one or more databases and click **Recover All Selected**.
5. Select the **Point-in-Time** check box.
6. From the **Restore Time** lists, select the date and the time.
7. Select the **Unconditionally overwrite existing database or files** check box.
8. Click **OK** to start the restore.

The database(s) will be restored to the selected date and time.

When you perform a point-in-time restore for a database, the next scheduled backup for that database will automatically convert to a Full backup.

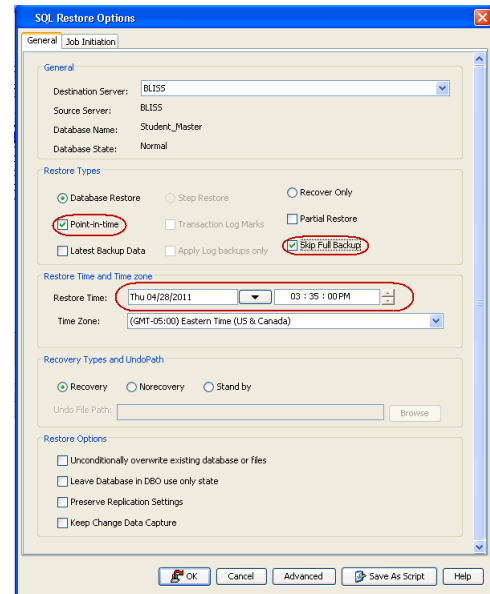


RESTORING ONLY THE LOGS

If you are using any third-party software to backup and restore the databases, you can use SQL Server iDataAgent to restore only the logs. The database for which you are restoring logs must be in the standby state. Follow the steps given below to restore logs for a database:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
2. Right-click the **<Instance>** and then click **All Tasks | Browse Backup Data**.
3. Click **OK**.
4. In the right pane of the Browse window, select a database that is already restored in the standby state and click **Recover All Selected**.
5. Select the **Destination Server**. Select the server and instance in which the restored database resides in the standby state.
6. Select the **Point-in-Time** check box.
7. Select the **Skip FULL Backup** check box.
8. Specify the **Restore Time**.
9. Click **OK** to start the log restore.

The most recent log backup before the **Restore Time** will be restored and applied to the database.



RESTORING SQL SERVER DATABASES TO A TRANSACTION MARK

You can restore directly to the marked transaction or a point before the transaction if the Transaction logs of the database has transaction marks.

This can be useful in situations where a transaction has caused undesirable results on a database and you need to go back to the point before that occurred.

Before performing the restore, ensure that transaction logs are backed up for the database that you want to restore to a transaction mark.

Follow the steps given below to restore a database(s) to a transaction mark:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
2. Right-click the **<Instance>** and then click **All Tasks | Browse Backup Data**.
3. Click **OK**.

- In the right pane of the Browse window, select a database that is already restored in the standby state.

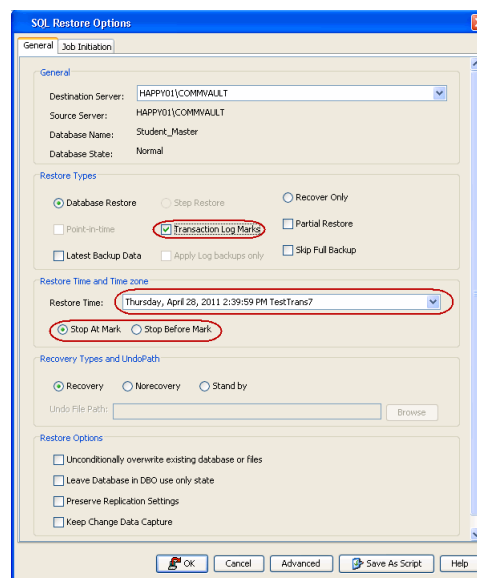
You can select multiple databases. However, the transaction marks must be common to all the selected databases.

- Click **Recover All Selected**.
- Select the **Transaction Marks** check box.
- Select the transaction mark from the **Restore Time** list.
- Select **Stop at Mark** or **Stop Before Mark**.

If you select **Stop at Mark** the database will be restored to the selected transaction mark. If you select **Stop Before Mark**, the database will be recovered to a previous transaction mark.

- Click **OK** to start the restore.

After performing a transaction mark restore job, the next backup job for that database is automatically converted to a full backup.



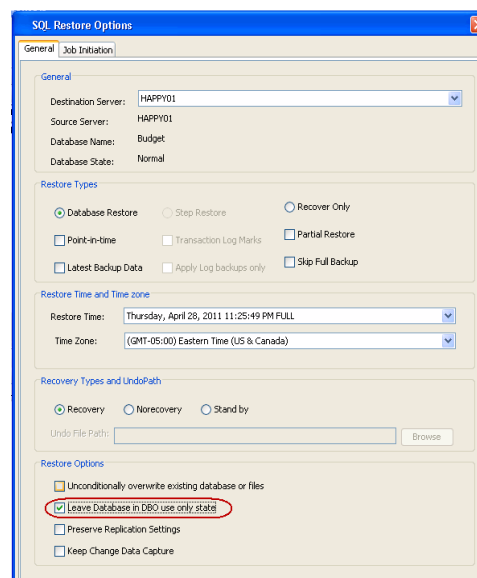
RESTORING A DATABASE WITH RESTRICTED ACCESS

This option is used to restore the database in a state where only the owner of the database can access the database.

Follow the steps given below to restore the database for the Database Owner (DBO_ONLY state):

- From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
- Right-click the **<Instance>** and then click **All Tasks | Browse Backup Data**.
- Click **OK**.
- In the right pane of the Browse window, select a database and click **Recover All Selected**.
- Select the **Leave Database in DBO use only state** check box.
- Select the **Unconditionally overwrite existing database or files** check box.
- Click **OK** to start the restore.

After the database is restored, only the owner of the database can access the database.



PRESERVING REPLICATION SETTINGS

You can preserve all the replication settings of a published database when you perform an out-of-place restore of a database. This options is used to restore backups of replicated databases. Only full backups are supported with this option, so make sure that you have performed at least one Full backup before restoring the database with replication settings. Follow the steps given below to preserve the replication settings:

- From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
- Right-click the **<Instance>** and then click **All Tasks | Browse Backup Data**.
- Click **OK**.
- In the right pane of the Browse window, select a database and click **Recover All Selected**.
- Select a SQL Server and instance from the **Destination Server** list.
- Ensure that **Recovery** option is selected from the **Recovery Types**.

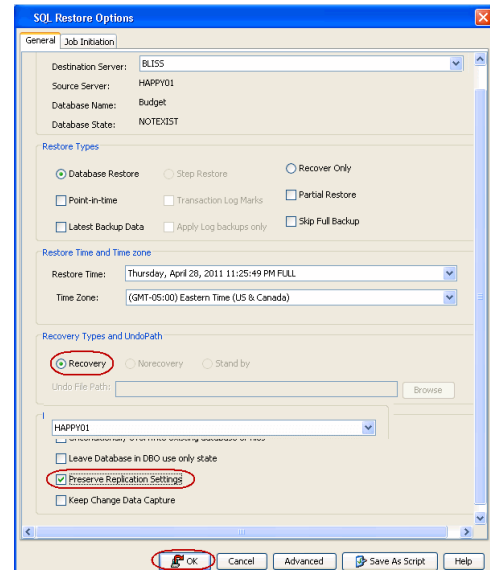
7. Select the **Preserve Replication Settings** check box.

Selection of Preserve Replication Settings option is not available in Microsoft SQL Server 2005.

8. Click **Advanced** and select the **Options** tab.
 - o Change the path of the database and log files under the **Physical Path** column.
 - o Click **OK**.
9. Click **OK** to start the restore.

After the restore has completed, all the replication settings of the database are available on the destination server.

If the publication database is restored, ensure that the replication settings of the `master` and `msdb` databases on the Distributor and Subscriber servers are consistent with the publication database. For more information, see [Backing Up and Restoring Replication Databases](#).



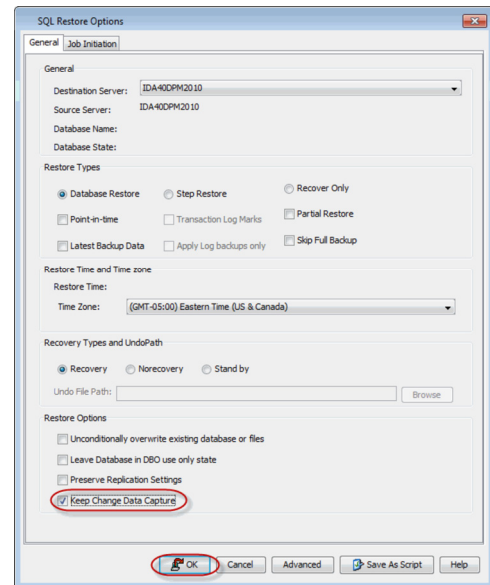
KEEPING CHANGE DATA CAPTURE (FOR SQL SERVER 2008 AND ABOVE)

The SQL Server 2008 has an option to capture the changed data. If you have enabled this option, the changed data (insert, update, and delete) is captured and applied to SQL tables. The details of the changes are available in an easily applicable relational format. You can restore the captured changes and apply them to the restored database. You cannot restore the captured changes if you are using VSS enabled backup for restore.

Follow the steps given below to restore the captured data:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
2. Right-click the **<Instance>** and then click **All Tasks | Browse Backup Data**.
3. Click **OK**.
4. In the right pane of the Browse window, select a database and click **Recover All Selected**.
5. Select the **Keep Change Data Capture** check box.
6. Click **OK** to start the restore.

After the restore has completed, all the data changes are available and applied to the restored database.



INCREASING THE VDI TIMEOUT

The VDI (Virtual Device Interface) is an API used to communicate with the SQL Server during all backup and restore operation. When restoring a database, the VDI timeout represents the time the system must wait for the SQL server to become ready to accept data into the database. If any databases are dropped from the SQL Server after the backup, the database file(s) do not exist on the sever. Before starting the restore operation, these files will be created automatically on the SQL Server. The time required to create the file(s) will be the same as when you first created the database or altered the database and added the file (s).

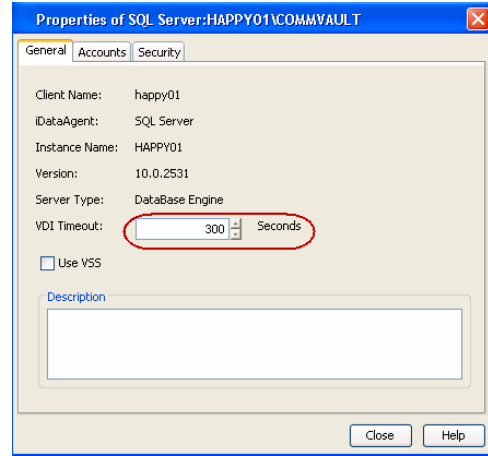
By default, the system will wait 50 minutes (10 retries of five minutes each) for the SQL server to respond to a VDI command before giving up. If the SQL server takes longer then this timeout value, the system assumes that the command failed, and it aborts the restore operation.

Follow the steps given below to increase the VDI Timeout for a SQL Server instance.

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL**

Server

2. Right-click the **<Instance>** and then click **Properties**.
3. Enter the **VDT Timeout**.
4. Click **OK**.



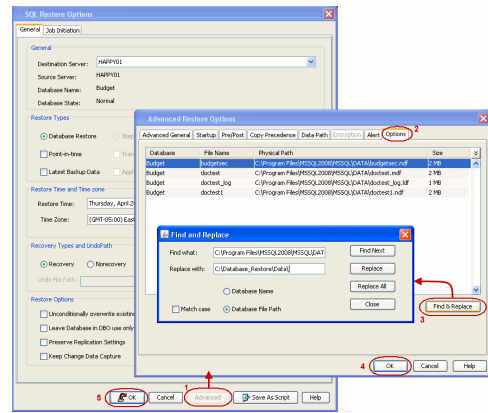
PERFORMING AN OUT-OF-PLACE RESTORE

Use one of the methods below to restore a SQL database to a different location.

COPYING A DATABASE TO A DIFFERENT LOCATION

Follow the steps below to restore a database to a different location on the same destination server:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
2. Right-click the **<Instance>** and then click **All Tasks | Browse Backup Data**.
3. Click **OK**.
4. In the right pane of the Browse window, select the databases you want to restore and click **Recover All Selected**.
5. Click **Advanced**.
6. Click the **Options** tab.
7. Rename the database under the **Database** column and change the path of the database and log files under the **Physical Path** column.
 - If there are many files, use the **Find & Replace** option to enter the new destination path for all of them at once.
 - If the database file path is changed without changing the database name, the existing database will be overwritten and will point to the new location.
8. Click **OK**.
9. Click **OK** to start the restore.



COPYING A DATABASE TO A DIFFERENT SQL SERVER INSTANCE

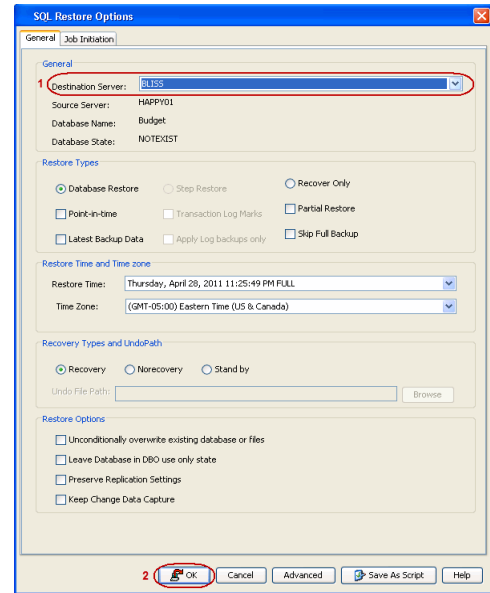
By default, when restoring to a different SQL Server instance, a database is restored to the same location in the destination server. Follow the steps below to restore a database to a different destination server.

System databases cannot be restored to a different server as they would overwrite the existing system databases in the destination server.

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
2. Right-click the **<Instance>** and then click **All Tasks | Browse Backup Data**.
3. Click **OK**.
4. In the right pane of the Browse window, select the databases you want to restore and click **Recover All Selected**.
5. Select the **Destination Server**.

If the same path cannot be used for any reason (e.g., the destination server does not have the corresponding path), click **Advanced**. From the **Options** tab, rename the database and change the path of the data files.

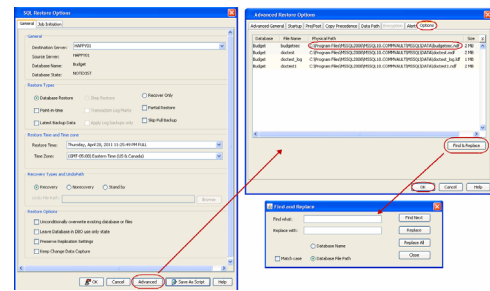
- Click **OK** to start the restore.



MOVING A DATABASE

You can restore the database to a different location without changing its name. You cannot move a system database to a different location.

- From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
- Right-click the **<Instance>** and then click **All Tasks | Browse Backup Data**.
- Click **OK**.
- In the right pane of the Browse window, select a database and click **Recover All Selected**.
- Click **Advanced**.
- Select the **Options** tab.
- Change the path of the database and log files under the **Physical Path** column.
- Click **OK**.
- Click **OK** to start the restore.



MOVING A DATABASE BACK TO THE ORIGINAL LOCATION AFTER A MOVE

If you perform a restore operation to move the database back to its original location, the database files will not be moved back as the database already exists on the server. The restore job will only check the database name and will restore the database on top of the same database regardless of its location. To move the database files back to the original location, use one of the following workarounds:

- Detach and re-attach the database.
- Perform a full backup of the database before moving the database back to its original location.
- Drop the database from the SQL Server instance and perform a restore with the default values.

COMMAND LINE RESTORES

You can perform restores of one of more databases from the command line interface.

Command line restores enable you to perform restore operations on multiple clients at the same time. It also allows you to reuse the command line scripts for additional restores.

When performing command line restores, note that backups taken from the CommCell Console can be restored using Command Line and vice versa. However, backups taken from a previous version of the CommCell Console can be restored only from the Command Line.

In order to run the restores from command line, you need an input xml file which contains the parameters for configuring the restore options. This input xml file can be obtained using one of the following ways:

- Download the input xml file template and save it on the computer from where the restore will be performed.
- Generate the input xml file from the CommCell Console and save it on the computer from where the restore will be performed.

LOG ON TO THE COMMSERVE

To run command line operations you must first login to the CommServe as follows:

- From Command prompt, navigate to <Software_Installation_Directory>/Base and run the following command:

```
qlogin -cs <commserve name> -u <user name>
```

- For example, to log on to CommServe 'server1' with username 'user1':

```
qlogin -cs server1 -u user1
```

PERFORM THE RESTORE

- Download the restore_template.xml file and save it on the computer from where the command will be executed.
- Execute the saved xml script using qoperation execute command.

```
qoperation execute -af restore_template.xml -clientName xxxxx -instanceName xxxxx -toTimeValue yyyy-mm-dd hh:mm:ss -
restoreSource xxxxx -database xxxxx
```

- Verify the status of the job using the following command:

```
qlist job -j JOBID
```

- Once the job completes, logout from the CommServe using the qlogout command.

```
qlogout [-cs commserver] [-all] [-tf tokenfile] [-tk token] [-h]
```

OUT-OF-PLACE RESTORE ON THE SAME SQL SERVER

- Download the out_of_Place_restore_template.xml file and save it on the computer from where the command will be executed.
- In the template locate the following section and edit it as per your setup:

```
<device>|DB1|#12!DB1_rename|#12!DB1|#12!E:\RestoreLocation\DB1.mdf|#12!C:\Program Files\Microsoft SQL
Server\MSSQL10_50.MSSQLSERVER\MSSQL\DATA\DB1.mdf</device>
```

```
<device>|DB1|#12!DB1_rename|#12!DB1_log|#12!E:\RestoreLocation\DB1_log.ldf|#12!C:\Program Files\Microsoft SQL
Server\MSSQL10_50.MSSQLSERVER\MSSQL\DATA\DB1_log.ldf</device>
```

Where:

- DB1 is the backed up database name.
 - DB1_Rename is the database name you want to provide for the restored database.
 - DB1 and DB1_log are the Logical Names of the backed up database.
 - E:\RestoreLocation is the destination path where you want to restore the database.
 - C:\Program Files\Microsoft SQL Server\MSSQL10_50.MSSQLSERVER\MSSQL\DATA is the source location of the backed up database.
 - #12! separates the values in the command. Note that the separator may vary based on the database, for e.g., #18!
- Execute the saved xml script using qoperation execute command.

```
qoperation execute -af out_of_Place_restore_template.xml -clientName xxxxx -instanceName xxxxx -restoreSource xxxxx -
database xxxxx -toTimeValue yyyy-mm-dd hh:mm:ss
```

- Verify the status of the job using the following command:

```
qlist job -j JOBID
```

- Once the job completes, logout from the CommServe using the qlogout command.

```
qlogout [-cs commserver] [-all] [-tf tokenfile] [-tk token] [-h]
```

OUT-OF-PLACE RESTORE ON A DIFFERENT SQL SERVER

- Download the out_of_Place_restore_template.xml file and save it on the computer from where the command will be executed.
- In the template locate the following section and edit it as per your setup:

```
<device>|DB1|#12!DB1_rename|#12!DB1|#12!E:\RestoreLocation\DB1.mdf|#12!C:\Program Files\Microsoft SQL
Server\MSSQL10_50.MSSQLSERVER\MSSQL\DATA\DB1.mdf</device>
```

```
<device>|DB1|#12!DB1_rename|#12!DB1_log|#12!E:\RestoreLocation\DB1_log.ldf|#12!C:\Program Files\Microsoft SQL
Server\MSSQL10_50.MSSQLSERVER\MSSQL\DATA\DB1_log.ldf</device>
```

Where:

- o DB1 is the backed up database name.
- o DB1_Rename is the database name you want to provide for the restored database.
- o DB1 and DB1_log are the Logical Names of the backed up database.
- o E:\RestoreLocation is the destination path where you want to restore the database.
- o C:\Program Files\Microsoft SQL Server\MSSQL10_50.MSSQLSERVER\MSSQL\DATA is the source location of the backed up database.
- o #12! separates the values in the command. Note that the separator may vary based on the database, for e.g., #18!

3. Execute the saved xml script using qoperation execute command, for example:

```
qoperation execute -af out_of_Place_restore_template.xml -clientName client1 -instanceName client1\instance1 -
restoreSource DB1 -database DB1 -toTimeValue yyyy-mm-dd hh:mm:ss -destClient/clientName client2 -
destinationInstance/clientName client2 -destinationInstance/instanceName client2\instance2
```

4. Verify the status of the job using the following command:

```
qlist job -j JOBID
```

5. Once the job completes, logout from the CommServe using the qllogout command.

```
qllogout [-cs commserver] [-all] [-tf tokenfile] [-tk token] [-h]
```

EXAMPLES

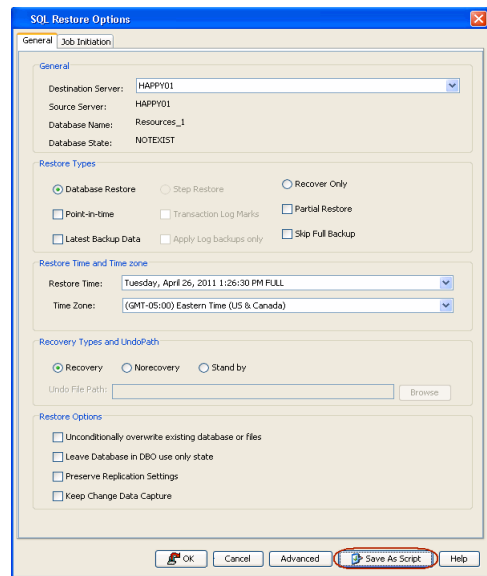
Restoring from a Current Backup	qoperation execute -af restore_template.xml -clientName client1 -instanceName client1 \instance1 -restoreSource DB1 -database DB2 -toTimeValue 2011-11-28 15:40:00
Restoring Databases to a Point in Time (date and time)	qoperation execute -af pit_restore_template.xml -clientName client1 -instanceName client1 \instance1 -restoreSource DB1 -database DB2 -toTimeValue 2011-11-28 15:40:00
Restoring Database In Place in Standby Mode	qoperation execute -af standby_inplace_restore_template.xml -clientName client1 -instanceName client1\instance1 -restoreSource DB1 -database DB2 -toTimeValue 2011-11-28 15:40:00 -sqlRecoveryType STATE_STANDBY -renameFilesSuffix C:\UndoPath

GENERATE THE COMMAND LINE SCRIPT FROM THE COMMCELL CONSOLE

In addition to the parameters provided in the template xml file, if you want to include additional options for the restore, you can do so by selecting the required options from the CommCell Console and generate the command line xml script for the restore operation.

Follow the steps given below to generate a script which you can use to perform a restore from the command line interface:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
2. Right-click the **<Instance>** and then click **All Tasks | Browse Backup Data**.
3. Click **OK**.
4. In the right pane of the Browse window, select the data that you want to restore and click **Recover All Selected**.
5. Select the required restore options which you want to execute using the script.
6. Click **Save as Script**.



7. Enter the location where you want to save the script or click **Browse** and navigate to the location.


The script will be saved as a .xml file and a .bat file is created.

If a file with the same name already exists in the specified location, the .xml file will be created with a timestamp. However, the .bat file will overwrite the existing file.

8. Enter the username and password for the user account which you want to use to perform the restore.

By default, the user account which you have used to login to CommCell console is used for performing the restore. However, if the user account does not have access to application and database, click **Use a different account**.

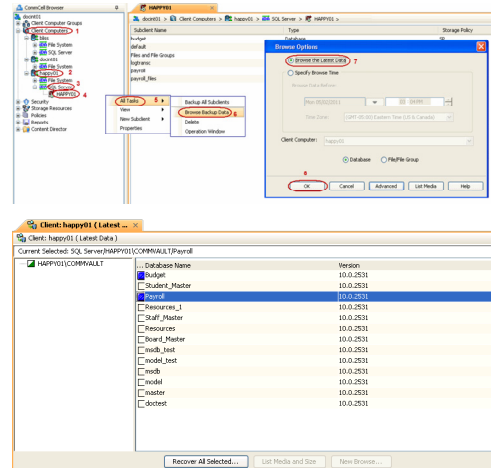
- Click **OK**.



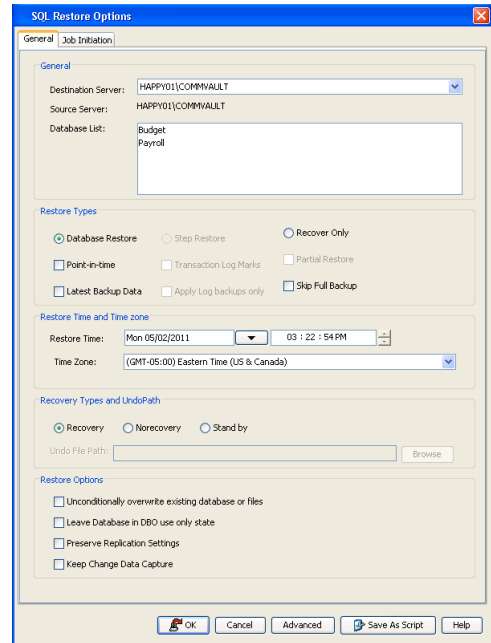
SCHEDULING A RESTORE

Follow the steps given below to schedule a restore job:

- From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
 - Right-click the instance and then click **All Tasks | Browse Backup Data**.
 - Click **OK**.
- Select a database to restore.
 - Click **Recover All Selected**.
- Select the required restore options. Click **OK**.

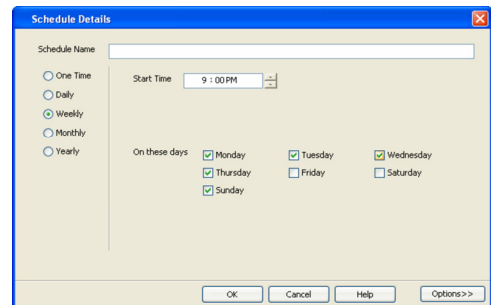
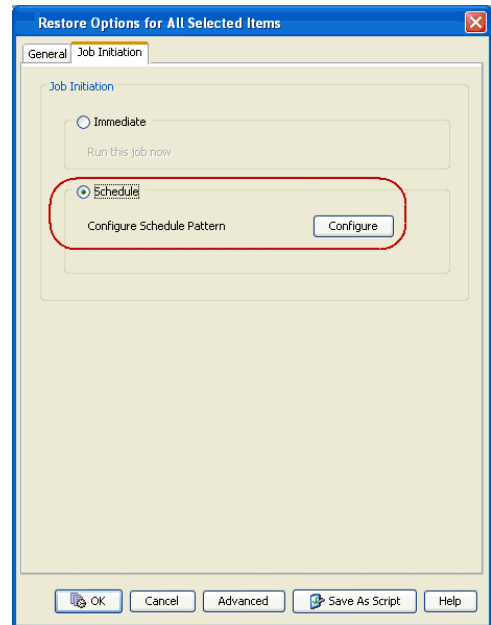


4.
 - Click **Job Initiation** tab and click **Schedule**.
 - Click **Configure**.



5. Select the appropriate scheduling options. For example:
 - Click **Weekly**.
 - Check the days you want the run the restore job.
 - Change the Start Time to 9:00 PM
 - Click **OK** to close the Schedule Details dialog
 - Click **OK** to close the Restore Options dialog

The restore job will execute as per the schedule.

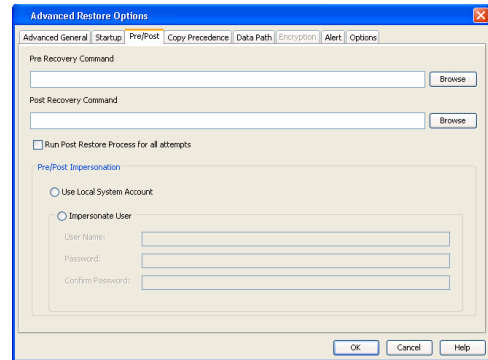


See Scheduling for a comprehensive information on scheduling jobs.

SETTING UP PRE-POST PROCESSES

You can run batch files or shell scripts before and/or after restore jobs. Follow the steps given below to setup a process before or after the restore job:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
2. Right-click the **<Instance>** and then click **All Tasks | Browse Backup Data**.
3. Click **OK**.
4. In the right pane of the Browse window, select a database and click **Recover All Selected**.
5. Click **Advanced** and then click the **Pre/Post** tab.
6. If you want perform a process before the restore job, enter the path for the batch file in the **Pre Recovery Command** box or click **Browse** to select the batch file.
7. If you want perform a process after the restore job, enter the path for the batch file in the **Post Recovery Command** box or click **Browse** to select the batch file.
8. Select one of the following options:
 - o **Use Local Accounts** - Select this option if the local account has permissions to execute the processes on the destination client.
 - o **Impersonate User** - Select this option and enter the username and password, that has the permissions to execute the processes on the destination client.
9. Click **OK**.
10. Click **OK** to start the restore.

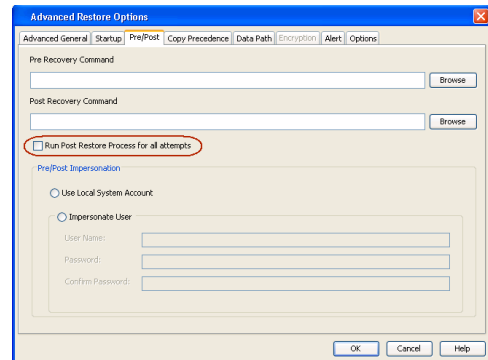


SETTING UP A PRE/POST PROCESS TO RUN AFTER EACH ATTEMPT

By default, a specified post process command is executed only on successful completion of the restore operation.

Use the following steps to run a post process even if the restore operation did not complete successfully. For example, this may be useful to bring a database online or release a snapshot.

1. From the CommCell Browser, navigate to **Client Computers | <Client> | SQL Server**.
2. Right-click the **<Instance>** and then click **All Tasks | Browse Backup Data**.
3. Click **OK**.
4. In the right pane of the Browse window, select a database and click **Recover All Selected**.
5. Click **Advanced** and then click the **Pre/Post** tab.
6. Enter the path for the batch file in the **Post Recovery Command** box or click **Browse** to select the batch file.
7. Select the **Run Post Process for all attempts** check box.
8. Click **OK**.
9. Click **OK** to start the restore.



MANAGING RESTORE JOBS

Once you initiate the restore operation, a restore job is generated in the Job Controller. Jobs can be managed in a number of ways. See Job Management for a comprehensive information on managing jobs.

The following sections provide information on the different job management options available:

RESTARTING JOBS

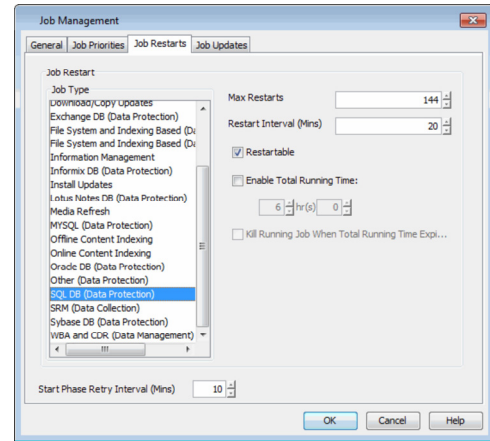
Jobs that fail to complete successfully are automatically restarted based on the job restartability configuration set in the Control Panel. Keep in mind that changes made to this configuration will affect all jobs in the entire CommCell.

To Configure the job restartability for a specific job, you can modify the retry settings for the job. This will override the setting in the Control Panel. It is also possible to override the default CommServe configuration for individual jobs by configuring retry settings when initiating the job. This configuration, however, will apply only to the specific job.

CONFIGURE JOB RESTARTABILITY AT THE COMMSERVE LEVEL

1. From the CommCell Browser, click **Control Panel** icon.
2. Select **Job Management**.
3. Click **Job Restarts** tab and select a **Job Type**.
 - o Select **Restartable** to make the job restartable.
 - o Change the value for **Max Restarts** to change the maximum number of times the

- Job Manager will try to restart a job.
 - Change the value for **Restart Interval (Mins)** to change the time interval between attempts for the Job Manager to restart the job.
4. Click **OK**.



CONTROLLING JOBS

The following controls are available for running jobs in the Job Controller window:

SUSPEND	Temporarily stops a job. A suspended job is not terminated; it can be restarted at a later time.
RESUME	Resumes a job and returns the status to Waiting, Pending, Queued, or Running. The status depends on the availability of resources, the state of the Operation Windows, or the Activity Control setting.
KILL	Terminates a job.

SUSPENDING A JOB

1. From the Job Controller of the CommCell Console, right-click the job and select **Suspend**.
2. The job status may change to **Suspend Pending** for a few moments while the operation completes. The job status then changes to **Suspended**.

RESUMING A JOB

1. From the Job Controller of the CommCell Console, right-click the job and select **Resume**.
2. As the Job Manager attempts to restart the job, the job status changes to **Waiting**, **Pending**, or **Running**.

KILLING A JOB

1. From the Job Controller of the CommCell Console, right-click the job and select **Kill**.
2. Click **Yes** when the confirmation prompt appears if you are sure you want to kill the job. The job status may change to **Kill Pending** for a few moments while the operation completes. Once completed, the job status will change to **Killed** and it will be removed from the Job Controller window after five minutes.

ADDITIONAL RESTORE OPTIONS

Several additional options are available to further refine your restore operations. The following table describes these options, as well as the steps to implement them.

Be sure to read the overview material referenced for each feature prior to using them.

OPTION	DESCRIPTION	RELATED TOPIC
Use hardware revert capability if available	<p>This option allow you to revert the data to the time when the snapshot was created. Selecting this option brings back the entire LUN to the point when the snapshot was created, overwriting all modifications to the data since the snapshot creation. This option is only available if the storage array used for SnapProtect Backup supports the revert operation.</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <Client> SQL Server. 2. Right-click the <Instance> and then click All Tasks Browse Backup Data. 3. In the Browse Options dialog box, click OK. 4. In the Client Browse window, select the database you want to restore and click Recover All Selected. 5. In the SQL Restore Options dialog box, click Advanced. 6. Select Use hardware revert capability if available. 7. Click OK. 	
Startup Options	The Startup Options are used by the Job Manager to set priority for resource allocation. This is	Refer to Job Priority and

	<p>useful to give higher priority to certain jobs. You can set the priority as follows:</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <Client> SQL Server. 2. Right-click the <Instance> and then click All Tasks Browse Backup Data. 3. In the Browse Options dialog box, click OK. 4. In the Client Browse window, select the database you want to restore and click Recover All Selected. 5. In the SQL Restore Options dialog box, click Advanced. 6. In the Advanced Restore Options dialog box, click Startup. 7. On the Startup tab, select Change Priority. 8. Enter a priority number - 0 is the highest priority and 999 is the lowest priority. 9. Select the Start up in suspended State check box to start the job in a suspended state. 10. Click OK. 	<p>Priority Precedence.</p>
<p>Copy Precedence</p>	<p>By default, the system retrieves data from the storage policy copy with the lowest copy precedence. If the data was pruned from the primary copy, the system automatically retrieves data from the other copies of the storage policy in the lowest copy precedence to highest copy precedence order. Once the data is found, it is retrieved, and no further copies are checked.</p> <p>You can retrieve data from a specific storage policy copy (Synchronous Copy or Selective Copy). If data does not exist in the specified copy, the data retrieve operation fails even if the data exists in another copy of the same storage policy. Follow the steps given below to retrieve the data from a specific storage policy copy:</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <Client> SQL Server. 2. Right-click the <Instance> and then click All Tasks Browse Backup Data. 3. In the Browse Options dialog box, click OK. 4. In the Client Browse window, select the database you want to restore and click Recover All Selected. 5. In the SQL Restore Options dialog box, click Advanced. 6. In the Advanced Restore Options dialog box, click Copy Precedence. 7. On the Copy Precedence tab, select the Restore from copy precedence check box. 8. Enter the copy precedence number. 9. Click OK. 	<p>Refer to Recovering Data From Copies.</p>
<p>Data Path Options</p>	<p>The data recovery operations use a default Library, MediaAgent, Drive Pool, and Drive as the Data Path. You can use this option to change the data path if the default data path is not available. Follow the steps given below to change the default data path:</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <Client> SQL Server. 2. Right-click the <Instance> and then click All Tasks Browse Backup Data. 3. In the Browse Options dialog box, click OK. 4. In the Client Browse window, select the database you want to restore and click Recover All Selected. 5. In the SQL Restore Options dialog box, click Advanced. 6. In the Advanced Restore Options dialog box, click Data Path. 7. On the Data Path tab, select the MediaAgent and Library. 8. Select the Drive Pool and Drive for optical and tape libraries. 9. Select the name of the Proxy server if you wish to restore using a proxy server. <p>You can select the Proxy server only when you are restoring from a snapshot. This option is not available for a regular restore.</p> <ol style="list-style-type: none"> 10. Click OK. 	<p>Refer to Change Data Path.</p>
<p>Encryption</p>	<p>If the client's data is encrypted with a pass phrase, you must enter the pass-phrase to start the data recovery operation. Follow the steps given below to enter the pass-phrase:</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <Client> SQL Server. 2. Right-click the <Instance> and then click All Tasks Browse Backup Data. 3. In the Browse Options dialog box, click OK. 4. In the Client Browse window, select the database you want to restore and click Recover All Selected. 5. In the SQL Restore Options dialog box, click Advanced. 6. In the Advanced Restore Options dialog box, click Encryption. 	<p>Refer to Data Encryption.</p>

	<ol style="list-style-type: none"> 7. Enter the Pass Phrase. 8. Click OK. 	
Alerts	<p>This option enables users or user groups to get automatic notification on the status of the data recovery job. Follow the steps given below to set up the criteria to raise notifications/alerts:</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers <Client> SQL Server. 2. Right-click the <Instance> and then click All Tasks Browse Backup Data. 3. In the Browse Options dialog box, click OK. 4. In the Client Browse window, select the database you want to restore and click Recover All Selected. 5. In the SQL Restore Options dialog box, click Advanced. 6. In the Advanced Restore Options dialog box, click Alerts. 7. Click Add Alert. 8. From the Add Alert Wizard dialog box, select the required threshold and notification criteria and click Next. 9. Select the required notification types and click Next. 10. Select the recipients and click Next. 11. Click Finish. 12. Click OK. 	Refer to Alert.
Command Line Restores	<p>Command Line Interface enables you to perform backups or restore from the command line. The commands can be executed from the command line or can be integrated into scripts.</p> <p>You can also generate command line scripts for specific operations from the CommCell Browser using the Save As Script option.</p>	Refer to Command Line Interface.
CommCell Readiness Report	<p>The CommCell Readiness Report provides you with vital information, such as connectivity and readiness of the Client, MediaAgent and CommServe. It is useful to run this report before performing the data protection or recovery job. Follow the steps given below to generate the report:</p> <ol style="list-style-type: none"> 1. From the Tools menu in the CommCell Console, click Reports. 2. Navigate to Reports CommServe CommCell Readiness. 3. Click the Client tab and click the Modify button. 4. In the Select Computers dialog box, clear the Include All Client Computers and All Client Groups check box. 5. Select the client from the Exclude list. 6. Click the Include > button. 7. Click OK. 8. Click the MediaAgent tab. 9. Clear the Include All MediaAgents checkbox. 10. Select the MediaAgent from the Exclude list. 11. Click Include >. 12. Click Run. <p>The generated report is displayed.</p>	Refer to CommCell Readiness Report.
Restore Job Summary Report	<p>The Restore Job Summary Report provides you with information about all the data recovery jobs that are run in last 24 hours for a specific client and agent. You can get information such as failure reason, failed objects, job options etc. It is useful to run this report after performing the restore. Follow the steps given below to generate the report:</p> <ol style="list-style-type: none"> 1. From the Tools menu in the CommCell Console, click Reports. 2. Navigate to Reports Jobs Job Summary. 3. Click Data Recovery on the General tab in the right pane. 4. On the Computers tab, select the client and the agent for which you want to run the report. 5. Click Run. 	Refer to Restore Job Summary Report.

Data Aging - SQL Server iDataAgent

TABLE OF CONTENTS

Getting Started

Extended Retention Rules

Data Aging for Transaction, Archive, and Logical Log Backups

Retention Rules for Log Backups

Data Aging for Stored Procedures

SQL Back in Time Restores and Data Aging Rules

Data Aging Rules for On Demand Backups

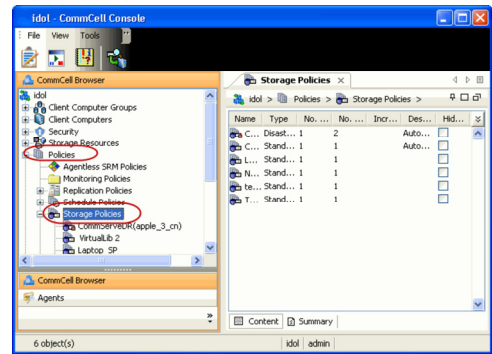
Enabling MSDB Database Clean-Up

GETTING STARTED

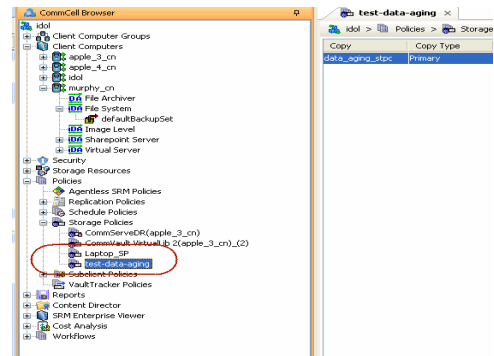
Data Aging is the process of removing old data from secondary storage to allow the associated media to be reused for future backups.

By default, all backup data is retained infinitely. However, you should change the retention of your data based on your needs. Note that if you continue to have infinite retention, you will also need infinite storage capacity.

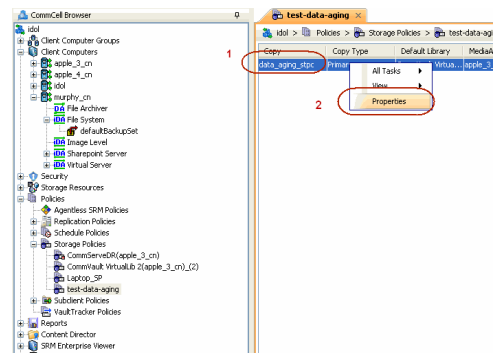
1. From the CommCell Browser, navigate to **Policies | Storage Policies**.




2. Highlight the **Storage Policy**.



3. From the right pane, right-click the **Storage Policy Copy** and click the **Properties**.




4.
 - Click the **Retention** tab.
 - Click the **Retain For** in the **Basic Retention Rules for All Backups** area.
 - Enter number of days to retain the data.
 - Enter number of cycles to retain the data.

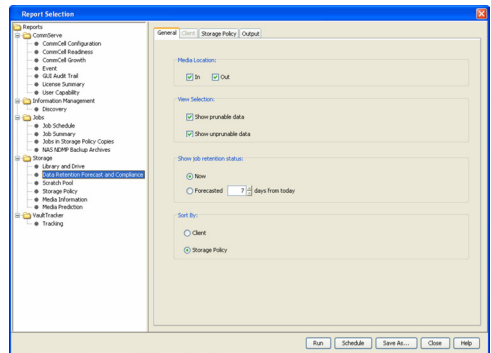


- Click **OK**.

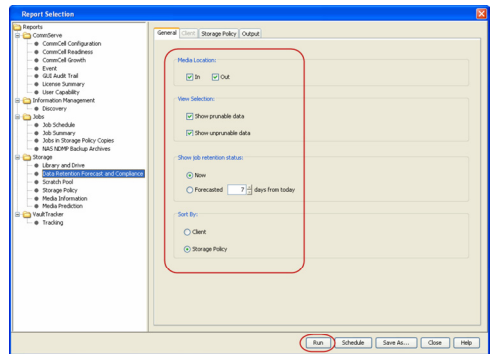
5. From the CommCell Browser, click the **Reports** icon.



6. Expand Reports and select **Data Retention Forecast and Compliance**.



7. Click **Run**.

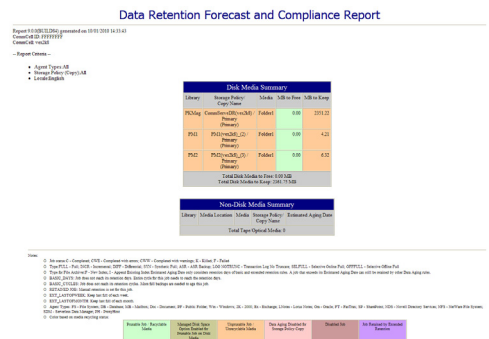


8. The report will display the data to be pruned when a data aging job is run.


To ensure only data intended for aging is actually aged, it is important to identify the data that will be aged based on the retention rules you have configured. Hence, ensure this report includes only the data you intend to age.

If necessary, fine-tune your rules so that only the intended data is aged.

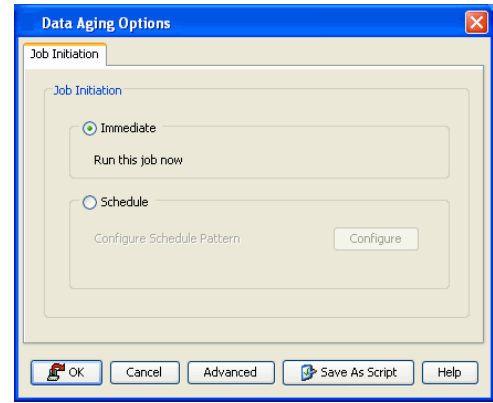
Once you run a data aging job, the data will be lost.



9. From the CommCell Console, right click the CommServe icon and click **All Tasks | Data Aging**.




10. Select **Immediate** in the Job Initiation section and click **OK**.



11. You can track the progress of the job from the **Job Controller** window. When the job has completed, the Job Controller displays **Completed**.

Make sure that the job completes successfully. If the job did not complete successfully, re-run the job.



EXTENDED RETENTION RULES

Extended retention rules allow you to keep specific full (or synthetic full) backups for longer periods of time.

Extended retention rules can be used in the following circumstances:

- If you have a single drive tape library
- If you want to create a hierarchical retention scheme (grandfather-father-son tape rotation)

In all other cases, it is recommended that the Auxiliary Copy feature be used for extended storage as it actually creates another physical copy of the data, thereby reducing the risk of data loss due to media failure.

UNDERSTANDING EXTENDED RETENTION RULES

Extended retention allows you to retain a specific full (or synthetic full) backup for an additional period of time. For example, you may want to retain your monthly full backups for 90 days.

Extended retention rules allow you to define three additional "extended" retention periods for full (or synthetic full) backups. For example:

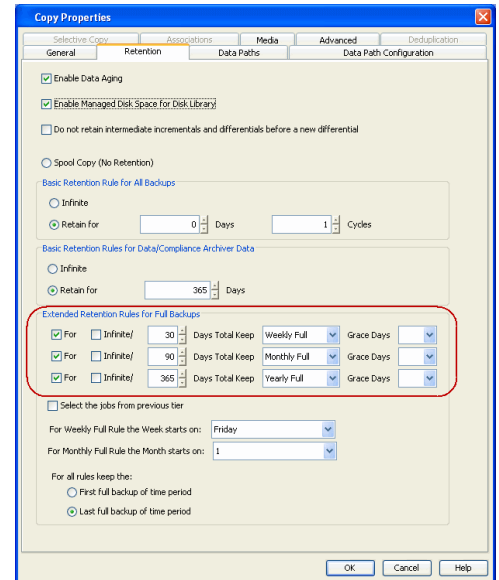
- You may want to retain your weekly full backups for 30 days.
- You may want to retain your monthly full backup for 90 days.
- You may want to retain your yearly full backup for 365 days.

A backup job will be selected for extended retention based on its start time. For example: If a backup job starts at 11:55 pm on August 31st and ends at 1 am on September 1st, then it will be selected as the last full backup for the month of August and will be picked up for extended retention.

SETTING UP EXTENDED RETENTION RULES

Use the following steps for setting up the extended retention rules:

1. Right-click the storage policy copy and click **Properties**.
2. Click the **Retention** tab.
3. Set the basic retention rules by clicking **Retain for** and entering the number of days and cycles appropriate for your organization.
4. Set the extended retention rules as follows:
 - Click the **For** button.
 - Enter the number of **Days Total** to retain the backup.
 - Click the **Keep** drop-down list, and select the desired backup criteria (e.g., Monthly Full).
 - Click the **Grace Days** drop-down list and select the number of days (e.g., 2).
5. Repeat Step 4 to configure additional extended retention.
6. Click **OK**.



DATA AGING FOR TRANSACTION, ARCHIVE, AND LOGICAL LOG BACKUPS

Log Backups (transaction, archive, or logical logs) are not considered part of the backup cycle. Therefore, storage policy cycle retention parameters do not apply to them.

RETENTION RULES FOR LOG BACKUPS

Log backups may be linked to data backup operations, which can allow storage policy cycle retention parameters to be applied to them.

This can be achieved as follows:

- If a full backup job is run on data, then the next log backup job will be linked to this full backup job.

These are considered as linked or chained log backups and are not aged until the linked data is aged. In addition, the following is also considered:

- Logs that need to be copied to secondary copies will not be aged both on primary and non-primary source copy
- Logs that exist only on one copy will be aged when they are older than the oldest data
- Logs that exist on multiple copies will be aged according to copy retention days
- Logs that exist on multiple copies with the longest retention days will be aged when they are older than the oldest data
- Partial, disabled logs will be aged when they are older than the oldest data
- If a full backup job is run on data and logs, then the next log backup will not be linked to this full backup job.

As this is an unlinked log backup, by default, this will follow the unique data aging rules for log backups. If you want such log backups to be aged according to the defined days retention rule for the data, you can do so as follows:

1. From the CommCell Browser, select **Tools | Control Panel**.
2. Double-click **Media Management**
3. Click the **Data Aging** tab.
4. Enable the **Prune All Database Agent Logs Only By Days Retention Rule** option.
5. Click **OK**.

LINKING FULL AND LOG BACKUPS FOR DATA RETENTION

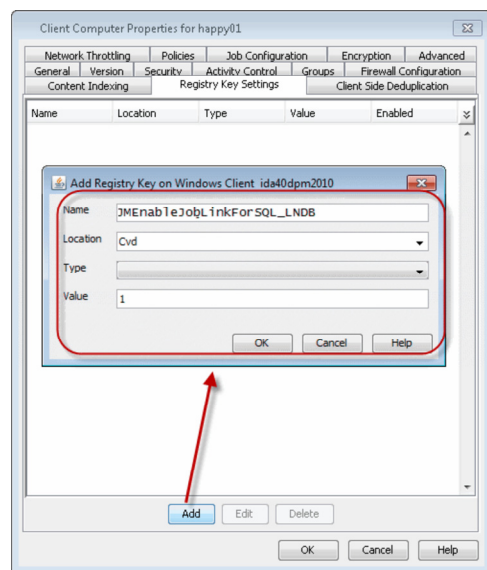
Log backups are linked to a full backup if they are run at the same time and the `JMEnableJobLinkForSQL_LNDB` key created in the client computer.

This is regardless of whether the full backup included data only or data and logs. Such backups follows the standard data aging rules.

You can create the key as follows:

1. From the CommCell Browser, 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 `JMEnableJobLinkForSQL_LNDB`.
6. In the **Location** list, select `cvd` from the dropdown list.
7. In the **Type** list, select `REG_DWORD`.
8. In the **Value** field, type any non zero value.
9. Click **OK**.



DATA AGING FOR STORED PROCEDURES

Data Aging for the SQL Server iDataAgents performs the following stored procedures that you may have been manually running on Enterprise Manager. When Data Aging is run, the system ages these histories from the CommServe database and the SQL Server.

- `sp_delete_backuphistory`
- `sp_delete_database_backuphistory`
- `sp_delete_backup_and_restore_history`

SQL BACK IN TIME RESTORES AND DATA AGING RULES

When you perform a back in time restore (i.e., restoring to a backup cycle earlier than the current backup cycle), all differential and transaction log backups which were run after the full backup from which the restored data was obtained will not be able to be aged until a new full backup is run. Running a full backup after performing a back in time restore releases the older backups and subsequent log backups for data aging.

DATA AGING RULES FOR BACKING UP ON-DEMAND SUBCLIENTS

Data Aging for On Demand backup jobs uses days/time, and ignores cycles and extended retention rules, as the determining factor for pruning the data. Therefore, once the retention time criteria has been met, all data (for both data and logs) is pruned that was backed up using the storage policy specified in the Command Line Interface.

An effective storage policy strategy for SQL On Demand backups is as follows:

- The same storage policy should not be used for regular backups and On Demand backups.
- The storage policy copy containing logs of On Demand backups should have a much longer retention time than other storage policies used by regular backups for the same instance. This is to prevent the logs of On Demand backups from being pruned before the data of regular backups, and allow the database to be fully restored and recovered using the data of old regular backups and logs afterwards.

ENABLING MSDB DATABASE CLEAN-UP

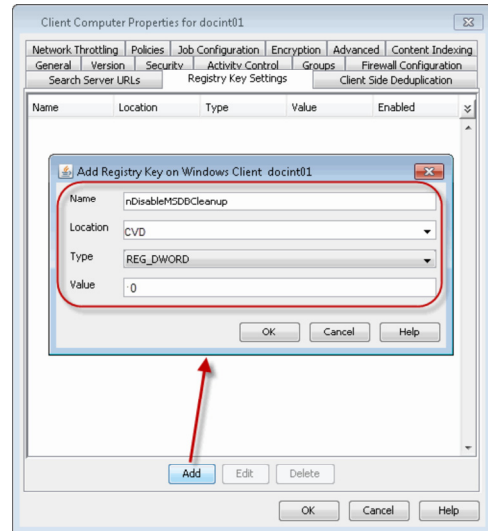
By default, Data Aging jobs do not perform a client-side clean-up of database metadata. However, to ensure that unnecessary data is not left behind, you can either use the system stored procedures mentioned below per SQL instance:

- `<sp_delete_backuphistory>`
- `<sp_delete_database_backuphistory>`
- `<sp_delete_backup_and_restore_history>`

Or enable client-side clean-up of database metadata process as follows:

1. From the CommCell Browser, navigate to **Client Computer**.
2. Right-click the **<CommServe Client>** and then click **Properties**.
3. Click the **Registry Key Settings** tab.
4. Click **Add**.

5. In the **Name** field, type nDisableMSDBCleanup.
6. In the **Location** list, type CVD.
7. In the **Type** list, select **REG_DWORD**.
8. In the **Value** field, type 0 to enable database clean-up.
9. Click **OK**.



ADVANCED TOPICS

Data Aging - Advanced

Provides comprehensive information on additional Data Aging capabilities.

Disaster Recovery - SQL Server iDataAgent

TABLE OF CONTENTS

Planning for a Disaster Recovery

Rebuilding the Operating System

Rebuilding the SQL Server Instance

Rebuilding the Master Database (Quick Rebuild)

Rebuilding Each SQL Server Instance (When Quick Rebuild Does Not work)

Restoring the SQL Server Instance

PLANNING FOR A DISASTER RECOVERY

Before you begin a disaster recovery, make sure to plan for disaster recovery as follows:

BUILDING A STANDBY SERVER FOR DISASTER RECOVERY

You can build a standby server for quick recovery and ensure that the standby database is always kept in a near ready state. If you plan to adopt this method follow the steps described in Replication Using Warm Database Restore.

MANAGING BACKUPS

As a best practice, it is recommended that you group databases into multiple subclients as follows:

- Add larger databases into separate subclients.
- Add small databases together into one or more subclients. This is important for the following reasons:
 - During Disaster recovery, when you have to quickly rebuild the entire instance, you can concurrently restore all the subclients together.
 - During Backup failures, the backup will restart from the beginning of the database instead of from the beginning of the entire instance. Similarly, this will ensure that large database backups are not affected by restarts from a smaller database.
- Once the subclients are created, schedule frequent backups for dynamic data and regular backup schedules for static data in the SQL Server.

REBUILDING THE OPERATING SYSTEM

In the case of disaster recovery, where a full system restore is required, you must first rebuild the system to exactly the state as it existed before the problem. Use the following steps to rebuild the operating system:

1. Rebuild the hardware if needed.
2. Install the same version of the Windows operating system with the same patches that were previously installed.
3. Configure the client with the same networking parameters and passwords that were previously set.
4. Edit the host file of the client to include the CommServe name.
5. Install the Windows File System iDataAgent on the client.

In a clustered environment, you must reinstall to the same virtual node as was being used previously.

6. Perform a full system restore of the Windows File System iDataAgent.

After you rebuild the operating system, if the SQL Server instance starts, use the steps in Restore the SQL Server Instance to restore the data.

REBUILDING THE SQL SERVER INSTANCE

After rebuilding the operating system, if the SQL Server Instance does not start, you need to rebuild the instance and then restore it to the current state. You can rebuild the instance by rebuilding the master database or by reinstalling each instance.

REBUILDING THE MASTER DATABASE (QUICK REBUILD)

The master database can be rebuild using SQL Server or SQL Server Express.

IF THE MASTER DATABASE EXIST

1. Insert the SQL Server installation media into the disc drive. If you are using SQL Server Express, download the Express kit and extract the contents to a local directory.
2. From the command prompt, navigate to the disc drive or the extracted directory and run the command for rebuilding the master database.

For SQL 2005:

```
start /wait setup.exe /qn INSTANCENAME=<InstanceName> REINSTALL=SQL_Engine REBUILDDATABASE=1 SAPWD=<NewStrongPassword>  
REINSTALLMODE=vomus
```

where:

/qn = suppresses all setup dialog boxes and error messages.

/qb = allows display of basic setup dialog boxes and error messages.

For SQL 2008:

```
Setup /QUIET /ACTION=REBUILDDATABASE /INSTANCENAME=InstanceName /SQLSYSADMINACCOUNTS=accounts [ /SAPWD=StrongPassword ]  
[ /SQLCOLLATION=CollationName ]
```

where:

/QUIET or /Q = setup runs without any user interface.

3. From the command prompt, run **Services.msc**
4. Right-click the SQL Server instance and select **Start**.
5. Restore the SQL Server Instance

IF THE MASTER DATABASE DOES NOT EXIST

In case when the master database is not available, use the following steps to rebuild the master database:

1. Navigate to the directory where the SQL Server is installed, usually C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL.
2. Open the Template Data folder and verify that there is a copy of the system databases created automatically during the original install.
3. Copy the missing files and paste them into the Data directory (C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\Data).
4. From the command prompt, run **Services.msc**
5. In the Services window, right-click the SQL Server instance and select **Start**.

REBUILDING EACH SQL SERVER INSTANCE (WHEN QUICK REBUILD DOES NOT WORK)

If the quick rebuild method does not work, you can rebuild the SQL Server by reinstalling each SQL Server instance.

1. Uninstall the existing SQL Server instances.
2. Install the SQL Server instances using the same installation paths and configuration settings as the previous installation.
Paths and instance name information of the original installation can be viewed by browsing the backup data from the CommCell Browser.
3. From the command prompt, run **Services.msc**
4. In the Services window, right-click each SQL Server instance and select **Start**.

RESTORING THE SQL SERVER INSTANCE

Restoring the SQL Server Instance involves restore of all the databases (system databases as well as the user created databases) of a SQL Server instance to the same computer.

Once you have rebuild the SQL Server instance, you can restore the instance using the following steps:

1. Stop any application or services that are accessing the SQL databases.
2. Restore the system databases.
3. Restore the user databases to a consistent point in time in the order that you require.

When restoring non-system databases, each database being restored runs as a separate job and a separate SQL process. Therefore, avoid starting more restores (jobs) than your SQL server and available memory can handle.

Additional Operations - SQL Server iDataAgent

TABLE OF CONTENTS

Audit Trail

Auxiliary Copy

License Administration

Online Help Links

Operating System and Application Upgrades

Operation Window

Schedule Policy

Storage Policy

Uninstalling Components

AUDIT TRAIL

The Audit Trail feature allows you to track the operations of users who have access to the CommCell. This capability is useful if a detrimental operation was performed in the CommCell and the source of that operation needs to be determined.

Audit Trail tracks operations according to four severity levels:

- **Critical:** This level records operations that will result in imminent loss of data.
- **High:** This level records operations that may result in loss of data.
- **Medium:** This level records changes to the general configuration of one or more entities. Such changes may produce unintended results when operations are performed.
- **Low:** This level records changes to status, addition of entities, and other operations that have minimal impact on existing CommCell functions.

To set Audit Trail retention periods:

1. From the **Tools** menu in the CommCell Console, click **Control Panel**, and then select **Audit Trail**.
2. From the **Audit Trail** dialog box, select the desired retention time (in days) for each severity level.
3. Click **OK**.

AUXILIARY COPY

An auxiliary copy operation allows you to create secondary copies of data associated with data protection operations, independent of the original copy.

1. Expand the **Policies** node, right-click storage policy for which you want to perform an auxiliary copy, click **All Tasks**, and then click **Run Auxiliary Copy**.
2. If you are starting the auxiliary copy operation from the CommServe level, select the storage policy for which you wish to perform the auxiliary copy.

If you are starting the auxiliary copy operation from the Storage Policy level, the Storage Policy field is already populated with the name of the Storage Policy you selected.

3. If the source copy is configured with a shared library, select the **Source MediaAgent** for the auxiliary copy.
4. Click **OK** to start the auxiliary copy operation. A progress bar displays the progress of the operation.

LICENSE ADMINISTRATION

REQUIRED LICENSES

The Microsoft SQL Server iDataAgent requires the **iDataAgent for Microsoft SQL Server** license.

CONVERTING EVALUATION LICENSES TO PERMANENT LICENSES

If you installed the software using an Evaluation License, you can upgrade to a Permanent License as follows:

1. From the CommCell Browser, right-click the CommServe icon, click **Control Panel**, and then click **License Administration**.
2. Select the **Update License** tab and then click **Convert**.
3. Check the box that corresponds to the evaluation license you would like to upgrade and then click **Convert**.

The license information is updated in the **License Administration** window.

USING CAPACITY-BASED LICENSING

License Usage by Capacity is a licensing mechanism that allows you to obtain licenses based on the amount of data you back up. It provides the following features:

- Flexibility of usage on all agents, rather than being tied to number of server, agents, etc.
- Allows you to purchase licenses based on your data protection needs

Both **Core** and **Enterprise** license types are available. Refer to License Usage by Capacity for comprehensive information on utilizing this method.

RELEASING A LICENSE

If you no longer require a license on a computer, such as cases where the computer is being retired, you can release the license and use it later for another computer. Backup data from the retired computer can still be restored after the license is released provided the data is not aged.

1. From the CommCell Browser, right-click the name of the client from which you want to release a license, click **All Tasks**, and then click **Release License for Client**.
2. Click **OK** to continue releasing the license
3. Click **Yes** to confirm you want to release the license or **No** to abort.

ONLINE HELP LINKS

Use the following links to view the online help for the corresponding tabs in the CommCell Console:

OPERATIONS	ENTITY	ONLINE HELP LINKS	SUB LINKS
CONFIGURATION	Agent	SQL Server Properties (General) SQL Server Properties (Version) SQL Server Properties (Security) SQL Server Properties (Activity Control) SQL Server Properties (Authentication)	
	Instance	Properties of SQL Server: <instance name> (General) Properties of SQL Server: <instance name> (Accounts) Properties of SQL Server: <instance name> (Security)	
	Subclient	Subclient Properties of <Subclient Name> (General) Subclient Properties of <Subclient Name> (Content) Subclient Properties of <Subclient Name> Backup Rules (Databases) Subclient Properties of <Subclient Name> (Backup Rules) (File/File Groups) Subclient Properties of <Subclient Name> (Pre/Post Process) Subclient Properties of <Subclient Name> (Storage Device) Subclient Properties of <Subclient Name> (Activity Control) Subclient Properties of <Subclient Name> (Encryption) Subclient Properties of <Subclient Name> (Security) Subclient Properties of <Subclient Name> (SQL Settings)	Database Configuration File / File Group Configuration
BACKUP	Backup Options	Backup Options for Subclient <Subclient Name>	Save As Script Command Line XML Options Schedule Details
	Advanced Backup Options	Advanced Backup Options (Data) Advanced Backup Options (Startup) Advanced Backup Options (Job Retry) Advanced Backup Options (Media) Advanced Backup Options (Data Path) Advanced Backup Options (VaultTracking) Advanced Backup Options (Alert)	
RESTORE	Restore Options	SQL Restore Options (General) SQL Restore Options (Job Initiation)	Save As Script Command Line XML Options Schedule Details

	Advanced Restore Options	Advanced Restore Options (Advanced General) Advanced Restore Options (Startup) Advanced Restore Options (Pre/Post) Advanced Restore Options (Copy Precedence) Advanced Restore Options (Data Path) Advanced Restore Options (Encryption) Advanced Restore Options (Alert) Advanced Restore Options (Options)	
--	--------------------------	---	--

OPERATING SYSTEM AND APPLICATION UPGRADES

Operating system upgrades are only supported when upgrading from one version of an OS to a different version of the same OS (e.g., Win2003 to Win2008). The two methods of upgrading are:

- **Seamless Upgrade** - This involves uninstalling the Agent software, upgrading the operating system, and then re-installing the Agent software.
- **Full OS Replacement** - This involves performing a clean install of a new version of the OS, re-installing any application software, then re-installing the CommServe, MediaAgent, and/or Agent software.

For Full OS Replacement, the client computer must be configured to have the CommServe, MediaAgent, and/or Client software re-installed to the same location, the same Fully Qualified Domain Name or short domain name, the same partitions, disk drive format (FAT, NTFS, et. al.), and IP configuration as previously.

If it is necessary to remove Agent software to facilitate an operating system or application upgrade, do not delete the icon for the Agent from the CommCell Console, or all associated backed up data will be lost.

Use the following strategy to upgrade the operating system software:

- Identify the computers you want to upgrade and the CommCell components installed on each of these computers.
- Choose the type of upgrade procedure you want to use on each computer: seamless or full replacement.
- CommServe, MediaAgent, and Client computers can be upgraded in any order.

OPERATION WINDOW

By default, all operations in the CommCell can run for 24 hours. To prevent certain operations from running during certain periods of the day, you can define operation rules so that these operations are disabled during those times.

When operation rules are configured, operations that are started within the time window specified will go to a queued (as opposed to pending) state. Once the time window specified in the operation rule has elapsed, these queued or running operations will resume automatically.

1. In the CommCell Browser, right-click the appropriate entity, click **All Tasks**, and then click **Operation Window**.
2. Click **Add**.
3. From the **Operation Window** dialog box:
 - Enter the name of the rule in the **Name** field.
 - Select either an administration, data protection (either full or non-full), and/or a data recovery operation from the **Operations** pane.
4. Click **OK**.

SCHEDULE POLICY

A schedule policy is a defined schedule or group of schedules for specific operations to be performed on associated objects within the CommCell. When the schedules from a policy are run, the specified operations, (e.g., auxiliary copy, backup, etc.) will be performed on the associated CommCell objects.

1. Expand the **Policies** node, right-click **Schedule Policies** and click **Add**.
2. Type the **Name** of the schedule policy.
3. Select the **Type** of schedule policy.
4. Select the **Agent Type**.
5. Type a description of the schedule policy.
6. Click **Add**.
7. Enter a Schedule Name in the **Schedule Pattern** tab.
8. Click **OK**.

9. On the **Associations** tab, select the objects to be associated with the schedule policy.
10. Click **OK**.

STORAGE POLICY

A Storage policy defines the data lifecycle management rules for protected data. Storage policies map data from its original location to a physical storage media and determine its retention period.

1. Expand the **Policies** node, right-click **Storage Policies**, and select **New Storage Policy**.
2. Click **Next**.
3. Select **Data Protection and Archiving** to create a regular storage policy or **CommServe Disaster Recovery Backup** to backup the CommServe database and click **Next**.
4. Click **Next**.
5. Enter the name of storage policy and click **Next**.
6. Enter the name of the primary copy and click **Next**.
7. From the drop down box, select the default library for the primary copy and click **Next**.
8. From the drop down box, select the MediaAgent and click **Next**.
9. Enter number of data streams and set the retention period for the policy and click **Next**.
10. Click **Next**.
11. Click **Browse**, browse to your designated deduplication store location and click **Next**.
12. Confirm your selections and click **Finish**.

UNINSTALLING COMPONENTS

You can uninstall the components using one of the following method:

- Method 1: Uninstall Components Using the CommCell Console
- Method 2: Uninstall Components from Add or Remove Programs

METHOD 1: UNINSTALL COMPONENTS USING THE COMMCELL CONSOLE

1. From the CommCell Browser, right-click the desired Client Computer and click **All Tasks -> Add/Remove Programs** and click **Uninstall Software**.
 2. **Uninstall Software Options** dialog will appear.
 3. In the **Uninstall Software** tab, select **Uninstall All** to uninstall all the software packages.
 4. In the **Job Initiation** tab, select **Immediate** to run the job immediately.
- You can track the progress of the job from the **Job Controller** or **Event Viewer**.

METHOD 2: UNINSTALL COMPONENTS FROM ADD OR REMOVE PROGRAMS

1. Click the **Start** button on the **Windows** task bar and then click **Control Panel**.
2. Double-click **Add/Remove Programs**.
For Windows Vista/Windows 2008, click Uninstall a Program in the **Control Panel**.
3. Click **Remove** to uninstall the components in the following sequence:
 1. <Agent>
 2. File System iDataAgent
 3. Base Software

ADVANCED TOPICS

Provides comprehensive information about additional capabilities for the following features:

- Audit Trail
- Auxiliary Copy
- License Administration
- License Usage by Capacity

- [Operation Window](#)
- [Schedule Policy](#)
- [Storage Policies](#)
- [Uninstalling Components](#)

[Back to Top](#)

Best Practices - SQL Server iDataAgent

CONCURRENTLY BACKING UP SQL SERVER WITH OTHER APPLICATIONS

When running backups using the Microsoft SQL Server iDataAgent, do not concurrently perform backups using other applications (e.g., Microsoft SQL Server Enterprise Edition). Doing so may compromise future restore operations.

There are a few jobs that SQL Server restricts during a backup. If one of these jobs are initiated while a backup is already in progress (or if a backup is initiated while one of these jobs is in progress) the backup job will terminate. These jobs are:

- creating or deleting database files
- creating an index
- performing any non-logged job
- shrinking the database.

CHANGING THE RECOVERY MODEL IN SQL SERVER

If you change the recovery model from Simple Recovery Model to either Bulk-Logged or to Full, it is recommended that you run a full database backup.

INCREASING THE SPEED OF SQL BACKUPS

If you want to increase the speed of SQL backups, you can increase the number of streams used by the backup to a number greater than 1 and enable software compression at the subclient level. Keep in mind that when you perform backups using multiple streams, you must have the same number of streams to restore the data.

You can set up multiple streams as described in [Configuring Data Streams](#).

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 a catch-all entity for client data. As a result, some data will not get backed up or scanned.

CREATING SUBCLIENTS

As a best practice, it is recommended that you add only a few small databases to each subclient and add larger databases into separate subclients. On a restart condition, the system will back up all databases in the subclient from the beginning. For this reason, you would not want a large database that has completed its backup successfully to be backed up again only because a smaller one has caused a restart to occur.

FILE AND FILE GROUP BACKUP AND RESTORE

If you plan to execute File/File Group restore jobs, plan to back up Transaction Logs regularly, because File/File Group restores cannot run without them.

Transaction Log Backups are required, refer to Microsoft article at <http://support.microsoft.com/kb/281122>.

BACKING UP AND RESTORING REPLICATION DATABASES

SQL Server replication allows distributing data to various servers throughout an organization.

It is recommended to backup the following replicated databases and their associated system databases regularly:

- The publication database at the Publisher
 - The master and msdb system databases at the Publisher and the publication database must be backed up at the same time.
- The distribution database at the Distributor
 - The master and msdb system databases at the Distributor and the distribution database must be backed up at the same time.
- The subscription database at each Subscriber
 - The master and msdb system databases at the Subscriber and the subscription database must be backed up at the same time.

Also make sure, when these databases are restored, the master and msdb database are consistent with the associated system database in terms of replication configuration and settings.

For backups and restores recommendations and strategies, refer to Microsoft article [Backing up and Restoring Replicated Databases](http://msdn.microsoft.com/en-us/library/ms151152.aspx) at <http://msdn.microsoft.com/en-us/library/ms151152.aspx>.

CHANGING STORAGE POLICIES

If you change a storage policy for a default subclient, and want to restore to a point in time prior to the change (i.e., restore from previous storage policy),

perform a Full backup of the database as soon as possible. Data from the old storage policy will not be considered as the most recent cycle and when it exceeds its retention period, and hence will be pruned.

SETTING PRIMARY REPLICA AS BACKUP PREFERENCE

It is recommended to set the primary replica as backup preference to ensure that complete backup cycles are available for restores later. Setting this preference will eliminate the possibility of the following errors:

1. **SKIPPING DATABASE [ALWAYSON_DATABASE]. NOT A PREFERRED REPLICA TO RUN BACKUP**

This error is displayed when the replica is not set as the backup preference and the SQL iDataAgent is set to recognize the backup preference setting of the availability groups. When a backup of the replica is attempted the databases which are a part of the availability group will be skipped.

2. **SKIPPING DATABASE [ALWAYSON_DATABASE]. ONLY BACKUP LOG IS FULLY SUPPORTED ON SECONDARY REPLICAS. BACKUP DATABASE SUPPORTS ONLY COPY-ONLY FULL BACKUPS OF THE DATABASE, FILES, OR FILE GROUPS. DIFFERENTIAL BACKUPS ARE NOT SUPPORTED ON SECONDARY REPLICAS**

This error will be displayed because the secondary replicas support copy only database backups. When backup of databases which are a part of availability group is attempted on secondary replica, such databases will be skipped from being backed up.

3. If transaction log backups are running on all replicas, and truncation happens during backup on secondary replica it will lead to breaking LSN chain on secondary replica.

Frequently Asked Questions - SQL Server iDataAgent

HOW DO I INSTALL SQL SERVER IDATAAGENT ON A COMPUTER WHERE MICROSOFT SQL SERVER SOFTWARE IS NOT INSTALLED?

Use the bIgnoreSQLServerCheck registry key to install the Microsoft SQL iDataAgent on a computer that does not have the Microsoft SQL server software installed.

IS MICROSOFT WINDOWS VISTA SUPPORTED?

Backups using the the SQL Server iDataAgent on Windows Vista are supported as long as the same User account (and not the local system account) is used to run the CommVault Communications Service (GxCVD) and for the SQL instance. See [Running Services Using a Windows User and User Accounts](#) for more details. [Back to Top](#)

HOW ARE TRANSACTION LOGS TRUNCATED?

Truncation of transaction logs use the following rules:

- Transaction logs will be truncated at the completion of a Transaction Log backup.
- Transaction logs will be truncated every time a checkpoint is processed, if the database is in truncate mode.
- Full and Differential backups do not truncate transaction logs.

CAN I RUN LOG AND FULL BACKUPS SIMULTANEOUSLY?

Yes. Full and log backup operations can run simultaneously for the same subclient. However, a full backup needs to exist prior to running them together. After a full has been run, subsequent parallel backups will work since the risk of breaking the chain will be lessened. When full and log backup operations are run simultaneously, the conversion rules will not be honored.

Note that running full and differential backups at the same time is not supported.

HOW DO I BACKUP A SQL DATABASE IN WINDOWS SHAREPOINT SERVICES (WSS)?

To successfully backup a SQL Database in a Windows SharePoint Services (WSS) environment, where the instance name of the Windows Internal Database resembles something like `MachineName\Microsoft##SSEE`, you need to:

1. Change the flag **Hide Instance** to **No** in **SQL Server Configuration Manager**→**Protocols for MachineName\Microsoft##SSEE**.
2. Ensure that the **SQL Server Browser** service is running in **SQL Server Configuration Manager**→**SQLServer 2005 Services**.
3. Ensure that **Named Pipes** is enabled in **SQL Server Configuration Manager**→**Protocols for MachineName\Microsoft##SSEE**.

WHEN DOES A NON-FULL BACKUP AUTOMATICALLY GET CONVERTED TO A FULL BACKUP?

A non-full backup is automatically converted to a full backup in the following situations:

- First backup of the subclient.
- Re-associating a subclient to another storage policy.
- Promote a secondary storage policy copy that is not synchronized with a primary copy (for all the subclients of a storage policy).
- If a backup job within the most recent backup cycle is pruned or disabled from a primary copy.
- Adding a new content path to the subclient.
- If you switch from a SnapProtect backup to a traditional backup or vice versa.
- CommCell Migration operation.
- After the following restore jobs the subsequent backup job for that database is automatically converted to a full backup:
 - Point in time restore
 - Transaction mark restore
 - Partial ("piecemeal") restore

WHAT HAPPENS IF I HAVE TWO DATABASES WITH THE SAME NAME?

If you set the SQL Server collation option to case-sensitive and two databases with the same name are created, the system will only discover one database.

WHAT DO I DO WHEN I UPGRADE THE SQL SERVER TO A NEWER VERSION

The system will automatically detect and display the newer version in the CommCell Console as follows:

- The SQL Server version information is refreshed and displayed in the instance **Properties** dialog box.

- The next backup that is run detects the new version.
- The database version of the backup is displayed during the browse operation.

However, if a SQL Server instance on which backups have been run is upgraded to a newer version and subsequently restored to a time prior to the SQL Server upgrade, the database will display the older version. To convert the database to a newer version, perform a full backup of the database after the restore.

HOW DO I RESTORE MULTIPLE SQL SERVER DATABASES TO A CONSISTENT TIME?

For some SQL server environments, you may have multiple databases for which absolute synchronization of the databases after a restore is critical. Transaction Mark restores and Point In Time restores are two means of keeping your databases synchronized.

- For transaction mark restores, you must administer and maintain your Microsoft SQL server transaction marks so that they are available for use by the SQL iDataAgent.
- Point in time restores use log backups to get your databases back to a specific minute and second. Backing up the tail of the log will be critical if you want to restore to just before the point of failure.

Follow the steps in Restoring SQL Server Databases to a Transaction Mark and Restoring SQL Databases to a Point in Time to restore the databases to a consistent time

HOW DOES TRANSPARENT BROWSE WORK FOR SQL SERVER?

The transparent inclusion during Browse operations provides the ability to restore data to a selected point in time. When restoring databases to a Transaction Mark or when using a Point in Time restores, the restore chain consists of the most recent full, the latest differential (if any), all subsequent log backups taken prior to the selected browse time, and, transparently, the next log backup after the browse time.

ARE THERE ANY CONSIDERATIONS WHEN RESTORING DATABASES BACKED UP BY VSS?

- VSS-enabled backup jobs must be restored to the original instance.
- For restore of databases backed up with VSS, databases can be moved to a new location on the server and can be renamed. However, data files cannot be renamed.
- To restore any system database that was previously backed up using VSS, the SQL Server services have to be stopped and restarted in single user mode.
- If the database is installed on the system drive, non-system databases will need to be manually brought back online after the restore.

WHAT ARE RESTORE CHAINS?

When you pick a restore time, the system determines the optimal combination of full, differential and log backups necessary to bring back the database(s) to the selected time. This combination of backups is the restore chain.

- For basic database restores, the restore chain consists of the most recent full, the latest differential (if any), and all subsequent log backups (if any) taken prior to the selected restore time.
- When restoring databases to a Transaction Mark or when using a Point in Time restore, the restore chain consists of the most recent full, the latest differential (if any), all subsequent log backups taken prior to the selected browse time, and, transparently, the next log backup after the browse time. The transparent inclusion enables restores to a selected point in time.

IN WHAT ORDER SHOULD SYSTEM DATABASES BE RESTORED?

- The master database must be restored first as it has to be online when other databases are restored.
- The msdb database should be restored second as it contains schedules.
- The model database can be restored third.

WHAT DOES THE "UNCONDITIONALLY OVERWRITE EXISTING DATABASE OR FILES" DO?

When this option is selected, the data being restored is unconditionally written to the specified location and overwrites the files of any database that is currently using the database name specified by you. This implements the T-SQL REPLACE command.

When this options is not selected, the restore job prevents the accidental overwrite of data files by failing the job.

WHAT IS THE DIFFERENCE BETWEEN INSTANCE RESTORE AND MULTIPLE DATABASES RESTORE?

When you perform an instance restore, all the databases are restored in one simultaneous job. However, if you select multiple databases that are not system databases, they are restored as multiple jobs.

HOW ARE THE TRANSACTION LOGS USED WHEN BROWSING THE BACKUP DATA FOR A FILE/FILE GROUP RESTORE?

When a browse operation is performed, files that existed as of the most recent full backup but were removed later will be included in the search results due to the replaying of the logs. Conversely, if a file was created after the most recent full backup, the search results will not present it for selection but the

transaction log will re-create this file as part the restore.

HOW CAN I RECLAIM SPACE AFTER TRANSACTION LOG IS TRUNCATED?

Use the DBCC SHRINKDATABASE command for Microsoft SQL Server to reduce the size of transaction logs, this will help reclaiming space. For using the command, refer to Microsoft article **DBCC SHRINKDATABASE (Transact-SQL)** at <http://msdn.microsoft.com/en-us/library/ms190488.aspx>.

Troubleshooting Backup - SQL Server iDataAgent

Backup | Restore

BACKUP FAILURES

The following section provides information on troubleshooting backups.

Database name contains [and] brackets	The use of embedded brackets '[' and ']' in database name may result in backup failures.
All data paths for the subclient are offline or busy	This error may be displayed if the Override Datapaths option is selected in the Data Paths dialog box in the Subclient Properties for a Log Storage Policy. This results in the Transaction Log backup operation waits for resources. To work around this issue, deselect the Override Datapaths option.
Time Out Failures	The default time allocated for backup and restore operations of SQL databases is 0 (infinite). If a backup or restore operation fails due to a timeout being reached, you can configure the nSqlQueryTimeout registry key to increase the amount of allocated time for backup or restore operations.
SQL Server jobs that cause backups to terminate	There are a few jobs that SQL Server restricts during a backup. If one of these jobs are initiated while a backup is already in progress (or if a backup is initiated while one of these jobs is in progress) the backup job will terminate. These jobs are: <ul style="list-style-type: none"> • creating or deleting database files • creating an index • performing any non-logged job • shrinking the database.
Backup chain is broken	When a full or differential backup is performed outside of the system, for example, from SQL Enterprise Manager, the subsequent log backups performed using SQL Server iDataAgent is set Do not convert log backups to full if log backup was performed using other software in the Subclient - Backup Rules tab. Make sure to enable the Disable Log Consistency Check in the Subclient - SQL Settings tab to ensure that the backup job completes successfully.

COMPLETED WITH ONE OR MORE ERRORS

Backup jobs from Microsoft SQL Server iDataAgent will be displayed as "Completed w/ one or more errors" in the Job History in the following cases:

- When a subclient which contains multiple databases is backed up, if one of the database is not backed up due to reasons like, database is in standby mode or database got corrupt etc, then job completes w/ one or more errors. The databases that failed will be shown as part of the failed items and those that were backed up will be shown as part of the successful items.
- When running a backup, a check is made to verify if the backup is restorable. If the log chain is broken the job will be displayed as Completed With Errors.
- A SQL backup job for a subclient with multiple databases will not retry backing up a single database if it fails. However, the job status will be displayed as Completed With Errors.

If the job goes into pending state, the job will restart from the point where it failed and if an attempt to back up the failed database has already been made, another attempt will not be performed.

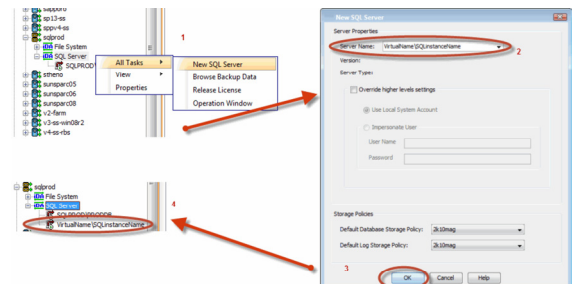
- For databases that are manually defined in a subclient but are inaccessible (e.g., it is not recognized, is offline, etc.), the job status for the backup will be displayed as Completed With Errors.

An event will be created for the inaccessible database during backup. If the inaccessible database is not needed, it can be permanently deleted from the subclient content.

VIRTUAL NODES ARE NOT BEING AUTOMATICALLY DISCOVERED IN A VERITAS CLUSTER ENVIRONMENT

Virtual nodes have to be configured manually after installing the SQL Server iDataAgent in a Veritas Cluster Environment from the CommCell Console. Follow the steps given below to configure the Virtual node:

1. From CommCell Console, navigate to the **SQL Server iDataAgent**.
2. Right-click the **SQL Server iDataAgent| All Task** and click **New SQL Server**.
3. Enter the SQL **Server Name** of the Virtual Node.
4. Click **OK**.
5. The virtual node will be visible under the selected **SQL Server iDataAgent**.



Troubleshooting Restore - SQL Server iDataAgent

Backup | **Restore**

BROWSE FAILURES

No Data to Restore	The following error message is displayed during a browse operation if CommServe and SQL Server time are not synchronized. There is no data to restore. Verify that the correct dates have been entered. Verify and ensure that the CommServe and SQL Server time are synchronized as close as possible. The SQL Server client can be ahead of the CommServe but not behind.
Time Out Failures	The default time allocated for backup and restore operations of SQL databases is 0 (infinite). If a backup or restore operation fails due to a timeout being reached, you can configure the nSqlQueryTimeout registry key to increase the amount of allocated time for backup or restore operations.
Time Out Failures for large databases	If the database you are attempting to restore is particularly large, it may become necessary to increase this time-out value. Re-run the restore after increasing the VDI Timeout value as described in Set the VDI Timeout Value for Large SQL Database Restores, see Increasing the VDI Timeout for step-by-step instructions.
Restore Failure of large Database	For large databases with over 2TB of Data, restore may fail due to not enough timeout in minutes. Set the timeout value of the query thread (for backup or restore) to a higher value, use nSQLQueryWaitMaxRetryCount registry key to change the timeout value.

COMPLETED WITH ONE OR MORE ERRORS

Restore jobs from Microsoft SQL Server iDataAgent will be displayed as "Completed w/ one or more errors" in the Job History in the following cases:

- When multiple databases are restored, if restore of one of the databases fail.

FILE SHARE RESOURCES NOT COMING BACK ONLINE

In MSSQL cluster, after restoring system database, the File Share Resources may not come backup online.

- Manually bring online the File Shared Resource using the Failover Cluster Manager.

RECOVERING DATA ASSOCIATED WITH DELETED CLIENTS AND STORAGE POLICIES

The following procedure describes the steps involved in recovering data associated with the following entities:

- Deleted Storage Policy
- Deleted Client, Agent, Backup Set or Instance

BEFORE YOU BEGIN

This procedure can be performed when the following are available:

- You have a Disaster Recovery Backup which contains information on the entity that you are trying to restore. For example, if you wish to recover a storage policy (and the data associated with the storage policy) that was accidentally deleted, you must have a copy of the disaster recovery backup which was performed before deleting the storage policy.
- Media containing the data you wish to recover is available and not overwritten.
- If a CommCell Migration license was available in the CommServe when the disaster recovery backup was performed, no additional licenses are required. If not, obtain the following licenses:
 - IP Address Change license
 - CommCell Migration license
 See License Administration for more details.
- A standby computer, which will be used temporarily to build a CommServe.

RECOVERING DELETED DATA

1. Locate the latest Disaster Recovery Backup which contains the information on the entity (Storage Policy, Client, Agent, Backup Set or Instance) that you are trying to restore.
 - You can check the Phase 1 destination for the DR Set or use Restore by Jobs for CommServe DR Data to restore the data.
 - If the job was pruned and you know the media containing the Disaster Recovery Backup, you can move the media in the **Overwrite Protect Media** Pool. See Accessing Aged Data for more information. You can then restore the appropriate DR Set associated with the job as described in Restore by Jobs for CommServe DR Data.

- o If the job is pruned and you do not know the media containing the Disaster Recovery Backup, you can do one of the following:

If you regularly run and have copies of the Data on Media and Aging Forecast report you can check them to see if the appropriate media is available.

2. On a standby computer, install the CommServe software. For more information on installing the CommServe, see CommServe Deployment.
3. Restore the CommServe database using the CommServe Disaster Recovery Tool from the Disaster Recovery Backup described in Step 1. (See Restore a Disaster Recovery Backup for step-by-step instructions.)
4. Verify and ensure that the **Bull Calypso Client Event Manager Bull Calypso Communications Service (EvMgrS)** is running.
5. If you did not have a CommCell Migration license available in the CommServe when the disaster recovery backup was performed, apply the IP Address Change license and the CommCell Migration license on the standby CommServe. See Activate Licenses for step-by-step instructions.
6. Export the data associated with the affected clients from the standby CommServe as described in Export Data from the Source CommCell.

When you start the Command Line Interface to capture data, use the name of the standby CommServe in the - `commcell` argument.

7. Import the exported data to the main CommServe as described in Import Data on the Destination CommCell.
This will bring back the entity in the CommServe database and the entity will now be visible in the CommCell Browser. (Press F5 to refresh the CommCell Browser if the entity is not displayed after a successful merge.)
8. If you have additional data that was backed up after the disaster recovery backup and before the deletion of the entity, use the procedure described in Import Metadata from a Tape or Optical Media to obtain the necessary information.
9. You can now browse and restore the data from the appropriate entity.

As a precaution, mark media (tape and optical media) associated with the source CommCell as READ ONLY before performing a data recovery operation in the destination CommCell.

SQL Server Properties (General)

Use this dialog box to manage the SQL Server properties.

Client Computer

Displays the name of the client computer or application server.

Billing Department

Displays the name of the billing department, when the Agent is associated with a billing department.

iDataAgent

Displays the identity of the agent that is installed on the client computer.

Installed

Displays the date on which the agent was installed or upgraded on the client computer or application server.

Auto Discover Instances

When selected, new instances are automatically discovered every 24 hours; or whenever the Communications Service (GxCVD) service is restarted.

By default this option is enabled, if it is selected during the Agent installation. (Express editions do not have the ability to select this option during the Agent installation, but is automatically selected after the installation.)

Manage Content Automatically

When selected, SQL databases that were automatically discovered but later deleted will be automatically deleted from the content of the default subclient when the next backup is run.

When cleared, databases that are manually added and later deleted are not automatically removed from the subclient when the next backup is run and the backup job will complete with errors.

By default, this is selected after client install or upgrade.

Description

Use this field to enter a description about the entity. This description can include information about the entity's content, cautionary notes, etc.

Version

Use this dialog box to view the iDataAgent version.

Security

Use this dialog box to:

- 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.

Available Groups

Displays the names of the user groups that are not associated with this CommCell object.

Associated Groups

Displays the names of user groups that are associated with this CommCell object.

Activity Control

Use this dialog box to enable or disable backups and restores on a selected subclient.

If data management and data recovery operations are disabled at the client computer group or client level, then these operations below these levels will be disabled. If data management / data recovery operations are enabled at the client computer group or client level, then these operations below these levels will be enabled.

Enable Backup

Specifies whether Backups will occur from this agent or subclient.

If cleared:

- Backup operations from this agent or subclient cannot be started and are skipped.
- Running and waiting data management operations for this agent or subclient run to completion.
- Stopped data management operations for this agent or subclient cannot be resumed until this option is enabled.
- Pending data management operations for this agent or subclient do not run until this option is enabled.

Enable Restore

Specifies whether Restores will occur from this agent or subclient.

If cleared:

- Restore operations from this agent or subclient cannot be started and are skipped.
- Running and waiting data recovery operations for this agent or subclient run to completion.
- Stopped data recovery operations for this agent or subclient cannot be resumed until this option is enabled.
- Pending data recovery operations for this agent or subclient do not run until this option is enabled.

Agent Properties (Authentication)

Use this dialog box to provide a valid user account to access to SQL Server to perform all operations, including backup, restore and browse.

Override Higher Levels Settings

Select to override the account settings configured at the Control Panel and Client Group levels for the selected SQL iDataAgent.

Use Local System Account

The Windows account configured to run the Communications Service (GxCVD) service and which is used by the system to perform all operations, including backup, restore and browse. By default, this is the Local System Account.

Impersonate User

Select to enter a user name and password for the Windows User Account that has permission to perform all operations, including backup, restore and browse. The account must already be set up on the client and must have Local administrator privileges and be a member of the SQL sysadmin fixed server role for the instance.

User Name

Use this space to type the name of the Windows user account that will have permission to execute the desired commands.

Password

Use this space to type the corresponding password for the account.

Properties of SQL Server: <instance name> (General)

Use this dialog box to manage the SQL Server Instance properties.

Client name

Displays the name of the client computer on which this instance is found.

Billing Department

Displays the name of the billing department, when the Instance is associated with a billing department..

iDataAgent

Displays the name of the iDataAgent to which this instance belongs.

Version

Displays the SQL version.

Server Type

Displays the server type according to the instance selected to be added.

VDI Timeout

Use the space to type the VDI Timeout value in seconds.

When restoring a SQL database, the VDI timeout represents the time the system must wait for the SQL server to become ready to accept data into the database. If the database you are attempting to restore is particularly large, it may become necessary to increase this time-out value.

Use VSS

Available for SQL Server clients running on Windows Server.

When selected, traditional full backups for all databases within the instance will be switched to VSS full backups.

When cleared, all full backups for all databases within the instance will use the traditional backup methods.

Description

Use this field to enter a description about the entity. This description can include information about the entity's content, cautionary notes, etc.

Subclient Properties of <subclient name>(General)

Use this dialog box to manage the SQL Server database properties.

Client Name

Displays the name of the Client computer to which this subclient belongs.

Billable Entity

Displays the name of the Billable Entity, when the subclient is associated with a billable entity.

Billing Department

Displays the name of the billing department, when the subclient is associated with a billing department.

iDataAgent

Displays the name of the iDataAgent to which this subclient belongs.

Instance

Displays the name of the Instance to which this subclient belongs.

Subclient

Displays the name of this Subclient.

You can use this space to enter or modify the name of the subclient.

Allow multiple data readers for Backup Copy

Specifies whether multiple data reads are allowed for a single Windows physical drive during backups on this subclient. This should be selected only for specialized hardware such as RAID, or possibly in the case of spanned volumes.

Disable Auto-Discovery

Specified to disable the auto-discovery and pre-assignment of SQL databases to the default subclient. By default, SQL databases are automatically discovered and assigned to the default subclient. (This option is available only in the default subclient.)

Description

Use this field to enter a description about the entity. This description can include information about the entity's content, cautionary notes, etc.

[Back to Top](#)

Subclient Properties of *<subclient name>* (Content)

Use this dialog box to define the contents of a new subclient or change the content of an existing subclient.

Database List (for databases)

Displays the names of databases that are currently included in the subclient content.

Database (for file/file groups)

Displays the names of database in which the file/file groups reside.

The names of file groups that are currently included in the subclient content are also displayed as a table.

Configure

Click to discover and add additional items to the subclient content.

Delete

Click to remove one or more items from the subclient content.

Properties of SQL Server: <subclient name> (Backup Rules)

Use this dialog box to enable or disable backup rules used to convert backup jobs under special circumstances to prevent failure conditions.

Convert

When selected, switches backup types as indicated under given circumstances.

When cleared:

- All backup jobs for master database will be converted to full backups.
- Backup jobs for msdb and model will run as is.
- For any user defined database, the system will fail the backup jobs that fits these given circumstances.

Log backups

- **Convert a log backup to full if a log backup was performed using other software**

Indicates that the next backup should be run as a full backup.

When a full or differential backup is performed outside of the system, for example, from SQL Enterprise Manager, the next log backups performed using SQL Server iDataAgent are not converted to a full backup. Therefore, it is recommended to run a full backup with the SQL Server iDataAgent.

- **Do not convert log backups to full if log backup was performed using other software**

Indicates that the next backup should be run as a log backup.

When using this option, make sure to enable **the Disable Log Consistency Check** in the **SQL Settings** tab to ensure that the backup job completes successfully. If this option is not selected, the backup job will fail with the error "Backup chain is broken".

File or File groups are added

When selected the next backup job is switched to a full backup.

When cleared, the selected backup type continues unchanged.

Properties of SQL Server: <instance name> (Backup Rules)

Use this dialog box to enable or disable backup rules used to convert backup jobs under special circumstances to prevent failure conditions.

Convert

When selected, switches file/file groups backup as indicated.

When cleared, the backup jobs will fail backup jobs in the given circumstances.

Pre/Post Process

Use this dialog box to add, modify or view Pre/Post processes for the selected subclient.

Pre Backup Process

Displays the name/path of the process that you want to run before the pre-backup phase.

You can use this space to enter a process that will execute before this phase, or use the **Browse** button to search for and select the name/path of the process. The system allows the use of spaces in the name/path, provided they begin with an opening quotation mark and end with a closing quotation mark.

Post Backup Process

Displays the name/path of the process that you want to run after the post backup phase.

You can use this space to enter a process that will execute before this phase, or use the **Browse** button to search for and select the name/path of the process. The system allows the use of spaces in the name/path, provided they begin with an opening quotation mark and end with a closing quotation mark.

Run Post Backup Process for all attempts

Specifies whether this process will execute for all attempts to run the phase.

When selected, this option will execute the specified process for all attempts to run the phase, including situations where the job phase is interrupted, suspended, or fails.

When cleared, the specified process will only execute for successful, killed, or failed jobs.

Run As / User Account

Displays either the Local System Account, or for added security, another account as having permission to run these commands.

Change

Click to add or modify the account that has permission to run these commands.

Storage Device

Use this dialog box to establish the storage device related settings on the selected subclient.

The following tabs are displayed:

- (Data) Storage Policy
- Data Transfer Option
- Deduplication

STORAGE POLICY

Use this tab to select or view storage policy settings on the selected subclient.

Data/Database/Transaction Log Storage Policy

Displays the storage policy to which this subclient is currently associated. To associate a storage policy to a new subclient or to change the storage policy associated with an existing subclient, click one in the list.

Incremental Storage Policy

Displays the name of the Incremental Storage Policy associated with this subclient, if the storage policy has the Incremental Storage Policy option enabled.

Data Paths

Click to view or modify the data paths associated with the primary storage policy copy of the selected storage policy.

Create Storage Policy

Click to launch the Create a Storage Policy wizard. Once the storage policy has been created, it will be displayed in the list of storage policies to which the selected subclient can be associated.

DATA TRANSFER OPTION

Use this tab to establish the options for data transfer.

Software Compression

Indicates whether software compression for the subclient or instance is enabled or disabled, and if enabled whether it is performed on the client or MediaAgent computer.

- **On Client**

Click to enable software compression on the client computer.

- **On MediaAgent**

Click to enable software compression on the MediaAgent computer.

- **Off**

Click to disable software compression.

Note that hardware compression has priority over the software compression. Hardware compression is established in the Data Path Properties dialog box. The above software compression option will take effect when the data path is associated with a disk library, or when hardware compression is disabled in the data path associated with tape libraries.



If the subclient is associated with a storage policy copy that is deduplicated, then the compression settings on the storage policy copy takes precedence. See Copy Properties (Deduplication) - Advanced tab for compression settings on deduplicated storage policy copy.

Resource Tuning

Indicates the processes used by the client to transfer data based and whether bandwidth throttling is enabled or not.

- **Network Agents**

Specifies the number of data pipes/processes that the client uses to transfer data over a network. Increasing this value may provide better throughput if the network and the network configuration in your environment can support it. On non-UNIX computers, the default value is 2 and a maximum of 4 can be established if necessary. On UNIX computers the default value is 1 and a maximum of 2 can be established if necessary.

- **Throttle Network Bandwidth (MB/HR)**

Specifies whether the backup throughput is controlled or not. (By default this option is not selected and therefore the throughput is not controlled). When selected, use the space to specify a value for the throughput. By default, this is set to 500. The minimum value is 1 and there is no limit to the maximum value. (In this case the backup throughput will be restricted to the maximum bandwidth on the network.)

Use this option to set the backup throughput, based on the network bandwidth in your environment. Use this option to reduce the backup throughput, so that the entire network bandwidth is not consumed, especially in slow links. Increasing this value will end up consuming the bandwidth with the maximum throughput limited to the network bandwidth capability.

Note that throttling is done on a per Network Agent basis.

DEDUPLICATION

Use this tab to establish the options for deduplication on the subclient. It indicates whether deduplication for the subclient is enabled or disabled, and if enabled whether the signature generation (a component of deduplication) is performed on the client or MediaAgent computer.

Note that deduplication is supported on disk storage devices. So the deduplication options are applicable only if the subclient is associated with a Storage Policy containing disk storage.

On Client

Click to enable signature generation on the client computer.

On MediaAgent

Click to enable signature generation on the MediaAgent computer.

[Back to Top](#)

Activity Control

Use this dialog box to enable or disable backups and restores on a selected subclient.

If data management and data recovery operations are disabled at the client computer group or client level, then these operations below these levels will be disabled. If data management / data recovery operations are enabled at the client computer group or client level, then these operations below these levels will be enabled.

Enable Backup

Specifies whether Backups will occur from this agent or subclient. .

If cleared:

- Backup operations from this agent or subclient cannot be started and are skipped.
- Running and waiting data management operations for this agent or subclient run to completion.
- Stopped data management operations for this agent or subclient cannot be resumed until this option is enabled.
- Pending data management operations for this agent or subclient do not run until this option is enabled.

Encryption

Use this dialog box to select the data encryption options for the selected content. When accessing this dialog box from the Subclient Properties Encryption tab, this setting applies only to the selected subclient content for operations run from the CommCell Console. When accessing this dialog box from the Instance Properties Encryption tab, this setting applies only to third-party Command Line operations. The functionality is not propagated to the Subclient Properties Encryption tabs.

None

When selected, no encryption will take place during a data protection operations.

Media Only (MediaAgent Side)

When selected, for data protection operations, data is transmitted without encryption and then encrypted prior to storage. During data recovery operations, data is decrypted by the client.

When using this setting in conjunction with the client property **With a Pass-Phrase**, you will be required to provide a pass-phrase for data recovery operations unless you export the client pass-phrase to the destination client(s). When using pass-phrase security for third-party Command Line operations or DataArchiver Agents stub recovery operations, you must export the pass- phrase to the destination client.

Network and Media (Agent Side)

When selected, for data protection operations, data is encrypted before transmission and is stored encrypted on the media. During data recovery operations, data is decrypted by the client.

When using this setting in conjunction with the client property **With a Pass-Phrase**, you will be required to provide a pass-phrase for data recovery operations unless you export the client pass-phrase to the destination clients.

Network Only (Agent Encrypts, MediaAgent Decrypts)

When selected, for data protection operations, data is encrypted for transmission and then decrypted prior to storage on the media. During data recovery operations, data is encrypted by the MediaAgent and then decrypted in the client.

When using this setting in conjunction with the client property **With a Pass-Phrase**, you will not be required to provide a pass-phrase for data recovery operations.

Script Preview

Click to display the backup script, based on the current subclient configuration, that will be submitted to RMAN when backups are performed for the selected Oracle subclient.

Properties of SQL Server: <subclient name> (SQL Settings)

Use this dialog box to specify tunable settings and log consistency checking for SQL backups.

Block Size

Specifies the block size that will be used during backup. All data transfers are in integral multiples of this value. Either accept the default value of 65,536 bytes (i.e., 64KB) or enter a value between 512 bytes and 65,536 bytes inclusive.

Buffer Count

Specifies the total number of buffers that will be used during backup. Either accept the default value of 1 or enter a value between 1 and INT_MAX (the maximum value of an int on the platform being used).

Note the following:

- The value cannot be less than the number of streams configured at the subclient level; otherwise, the number of streams will be used during backup.
- If the buffer count is higher than the streams, the buffer count is used during backup.

Maximum Transfer Size

Specifies the maximum transfer size that will be used between during backup. Either accept the default value of 65,536 bytes (i.e., 64KB) or enter a value in multiples of 64 KB. The range is between 65,536 bytes and 4,294,967,296 bytes (i.e., 4 MB).

Disable Log Consistency Check

By default, this option is enabled and the software will check for log consistency during backup. If detected that the backup chain is broken when this option is enabled, the job will not proceed. Select this option to disable log consistency checking. When disabled, if the software detects that the chain is broken, the job will not fail but the database being backed up may not be restorable later on.

Database Configuration

Use this dialog box to discover and change the databases included in the subclient.

Database Name

Lists the names of the databases for inclusion in the subclient.

Subclient Name

Displays the subclient to which the database is assigned. Click this space to assign the database to one of the following:

- Default subclient
- Any user-defined subclient
- Do Not Backup subclient

Change all selected databases to

When more than one database is selected, reassigns the selected databases to the selected subclient.

Discover

Discovers new databases not yet included in the subclient contents.

File/File Group Configuration

Use this dialog box to discover and change the files /file groups included in the subclient. .

Database

Displays the name of the database whose current file groups and files are displayed in the table below. To change the database name, click one in the list. The following information is displayed for the files and file groups that have been discovered through clicking **Discover**.

- **Name**

Displays the names of the file groups and files for the database.

- **Type**

Displays the type of data (file group or file) within the database

- **Physical File Name**

Displays the directory path for the included files.

- **Subclient**

Displays the subclient to which the database is assigned. Click this space to assign the database to an user-defined subclient to the database.

Discover

Discovers any file groups/files not yet configured within the system, allowing them to be assigned to subclients other than the default.

Change all selected FFG to

When more than one file group/file is selected, reassigns the selected file groups/files to the selected subclient.

Backup/Archive Options

Use this dialog box to schedule or immediately run a backup/archive job. Note that all the options described in this help may not be available and only the options displayed in the dialog box are applicable to the agent for which the information is being displayed.

Select Database Backup Type

- **Full**

Specifies the job as a Full backup, which backs up all data for the selected subclient(s).

- **Differential**

Specifies the job as a Differential backup, which backs up only that portion of the data that is new or has changed since the last full backup.

- **Transaction Log**

Specifies the job as a transaction log backup. This will back up the transaction log, providing point-in-time restore functionality.

- **Do not truncate log**

Specifies the operation will back up the transaction log in situations where the database is damaged or has not been recovered.

Job Initiation

- **Run Immediately**

Specifies this job will run immediately.

- **Schedule**

Specifies this job will be scheduled. Click **Configure** to specify the schedule details.

Configure Alert

- **Alert**

The currently configured Alert.

- **Add/Modify Alert**

When clicked, opens the Alert Wizard to configure alerts for this operation.

- **Delete Alert**

When clicked, deletes any existing alerts that are already configured.

Advanced

Click to select advanced backup/archive options, such as Start Log After Successful Backup, Partial (Excludes read-only Filegroups), and Tail-log Backup.

Save As Script

Click to open the Save As Script dialog, which allows you to save this operation and the selected options as a script file (in XML format). The script can later be executed from the Command Line Interface using `qoperation execute` command.

When you save an operation as a script, each option in the dialog will have a corresponding xml parameter in the script file. When executing the script, you can modify the value for any of these XML parameters as per need.

To view the XML values for each of the options in the dialog, see the following:

- Command Line XML Options for Microsoft SQL Server iDataAgent
- Command Line XML Options for MySQL iDataAgent

[Back to Top](#)

Save as Script

Use this dialog box to choose a name and path for the script file and the mode of execution.

Client

Enter or select the name of the client computer where the script will be created.

Path

Enter the path for the script that will be created.

Browse

Click this button to browse to a path for the command line script.

- Scripts are not supported on the Windows NT platform.
- It is recommended not to use any reserved device names (e.g., LPT1) as the name of the file.
- The file names are not case-sensitive.
- Do not end the file name with a trailing space or a period. Although the underlying file system may support such names, the operating system does not support them.

Mode

• Synchronous

Specifies that the script execute in synchronous mode. A synchronous operation exits only when the operation has completed. This option is only available when scripting a single job.

• Asynchronous

Specifies that the script execute in asynchronous mode. An asynchronous operation submits the job to the CommServe and exits immediately, returning control to the calling program or script.

Specify User Account to Run the Script

Specifies to use the given user account to save the operation as a script.

• Use the Currently Logged in User Account

Click to use the same user account used for logging into the CommCell Console.

○ User Name

Type the user name that was used for logging into the CommCell Console.

○ Password

Type the password for the user account used for logging into the CommCell Console.

○ Confirm Password

Type to re-confirm the password.

• Use a Different User Account

Click to specify a different user account to save the operation as a script.

○ User Name

Enter the different user name to be used for saving the operation as a script.

○ Password

Enter the password for the user account to be used for saving the operation as a script.

• Confirm Password

Type to re-confirm the password.

[Back to Top](#)

Schedule Pattern

Use this dialog box to schedule jobs.

Schedule Name

Displays the name of the schedule. If creating a new schedule, use this space to enter the name of the schedule.

One Time

For a job run a single time. Select the date and time you want the job to begin.

Daily

For a job run on a daily basis. Choose the interval, in days, at which you want the job to repeat.

Weekly

For a job run on a weekly basis. You can run the job every n number of weeks on the selected days of the week.

Monthly

For a job run on a monthly basis. You can run the job every n number of months on the selected date or day. You can also select the standard calendar or a custom calendar (if available). Monthly schedules using custom calendars will run according to the definitions of a month as defined in the calendar.

Yearly

For a job run on a yearly basis. You can run the job annually on the selected date or day.

The Monthly and Yearly selections allow you to schedule other calendar events that are commonly of interest. For example, you can schedule backups on the last weekday every three months for quarterly backups. Optionally, you can select the **Every n Month(s)** option, which allows you to customize intervals at which the job will run. You can also select the standard calendar or a custom calendar (if available). Yearly schedules using custom calendars will run according to the definitions of a year as defined in the calendar.

Automatic Schedule

Select this option to run a job on a specified frequency.

Job Interval

Allows you to specify intervals for jobs.

- **Minimum Interval between Jobs**

Specify the number of hours to start a backup job if the below criteria (Network Management, Power Management and Resource Utilization) is/are satisfied.

- **Maximum Interval between Jobs**

Specify the number of hours since the last successful job completed to start a job even if the below criteria (Network Management, Power Management and Resource Utilization) is/are not satisfied.

Network Management

Allows you to specify the available network types.

- **Start only if wired network is available**

Select to start the job when the network connectivity is wired.

- **Start only if specific network is available**

Select to start a job using a specific network. Specify network details and add the network address.

- **Minimum Network Bandwidth n Kbps**

Select to specify the minimum number of kbps that the job should use for the network bandwidth. By default, the minimum bandwidth is 128.

- **Throttle at n % until bandwidth n Kbps**

Select to specify the percentage of throttle that the job should use until the bandwidth reaches an specific kbps value. By default, the job throttles at 40% until the bandwidth is 2048.

Power Management

Allows you to specify power management options for the computer.

- **Start only if the computer is running on A/C power**

Select this option to run the job only when the computer is on A/C power.

- **Stop the task if batter mode begins**

Select this option to stop the job if the computer is on battery mode.

Resource Utilization

Allows you to setup the job schedule with specific resource utilization.

- **Start only if CPU utilization is below**

Select this option to run the job when CPU utilization is below the specified percentage.

- **Start the job if free disk space drops below**

Select this option to run the job when disk space is below the specified percentage.

Options

Click to display **Advanced Scheduling Options**.

Advanced Backup/Archive Options (Data)


You can select advanced backup/archive data options for the operation. Note that all the options described in this help may not be available and only the options displayed in the dialog box are applicable to the agent or enabler for which the information is being displayed.

Start Log Backup after Successful Backup

Specifies that a Transaction Log backup will start automatically after a successful Full or Differential backup operation is completed. This is useful when you want to back up logs immediately after a data backup, and allows you to do so without creating two scheduled jobs.

SQL Native Backup Compression

For SQL Server 2008 Enterprise or later, specifies that the backup will be compressed by SQL Server before being backed up. The size of the backup will be smaller than an uncompressed backup of the same data. Typically, compressing a backup will require less device I/O which should increase backup speed significantly. However, CPU usage may increase for compressed backups and you may want to evaluate performance counters. Scheduling the backup during off-peak hours or compressing only low-priority backups may be desirable. When using the SQL Native Backup Compression option, there is no need for deduplication as the data will be compressed with this option and deduplication will not consequently save any more space.

	VSS-enabled and SnapProtect backups are not supported backup types for compression.
---	---

Partial (Excludes read-only Filegroups)

A partial backup resembles a full database backup, but does not contain all the filegroups. Instead, it contains all the data in the primary filegroup, every read/write filegroup, and any optionally-specified read-only files. Partial backups are useful whenever you want to exclude read-only filegroups. A partial backup of a read-only database contains only the primary filegroup.

Tail-log Backup (database is left in restoring state)

This advanced backup option is only available if the **Transaction Log** option was selected in the Backup Option dialog box. Specifies to back up the tail of the log and to leave the database in the RESTORING state.

Copy Only

A copy-only backup is an independent backup and does not affect the usual chaining involved in traditional backups. This type of backup can be performed for a special purpose and does not affect the overall backup and restore procedures. It applies to full backups, full Volume Shadow Services (VSS)-enabled backups, and full SnapProtect backups.

After taking a full copy-only backup, log backups cannot be run unless a traditional full backup is present. Differential backups can be run but they cannot be restored unless a traditional full backup exists prior to the differential.

Create Backup Copy immediately

Select to create an inline backup copy to start movement of snapshot to media, immediately after the completion of the SnapProtect backup job.

Startup

Select from the following options. Note that all the options described in this help may not be available and only the options displayed in the dialog box are applicable to the operation for which the information is being displayed.

Priority

- **Use Default Priority**

If selected, the default priority for this type of job will be used in determining how the Job Manager will allocate resources for this job.

- **Change Priority**

Use this option to manually specify the priority for the job, between 0 (highest priority) and 999 (lowest priority). The Job Manager will use the priority setting when allocating the required resources. This is useful if you have jobs that are very important and must complete, or jobs that can be moved to a lower priority.

Start up in suspended state

Specifies that this job will start in the Job Controller in a suspended state and cannot run until the job is manually resumed using the **Resume** option. This option can be used to add a level of manual control when a job is started. For example, you could schedule jobs to start in the suspended state and then choose which scheduled jobs complete by resuming the operation started in the suspended state.

Description

Use this field to enter a description about the entity. This description can include information about the entity's content, cautionary notes, etc.

Job Retry

Note that all the options described in this help may not be available and only the options displayed in the dialog box are applicable to the agent or operation for which the information is being displayed.

Enable Total Running Time

The maximum elapsed time, in hours and minutes, from the time that the job is created. When the specified maximum elapsed time is reached, as long as the job is in the "Running" state, it will continue; if the job is not in the "Running" state when the specified time is reached, Job Manager will kill the job.

Enable Number of Retries

The number of times that Job Manager will attempt to restart the job. Once the maximum number of retry attempts has been reached, if the job has still not restarted successfully, Job Manager will kill the job. Note that this job-based setting will not be valid if restartability has been turned off in the Job Management Control Panel.

Kill Running Jobs When Total Running Time Expires

Option to kill the job when the specified Total Running Time has elapsed, even if its state is "Running". This option is available only if you have specified a Total Running Time.

Advanced Backup Options (Media)

You can select advanced backup media options for the operation.

Start New Media

This option starts the backup/migration/archive operation on a new media, which causes the following to occur:

- If removable media is used, the current active media is marked as Appendable and a new media is used for the backup/migration/archive.
- If disk media is used, a new volume folder is created for the backup/migration/archive.

If cleared, the operation automatically uses the current active media or volume.

Mark Media Full after Successful Operation

This option marks media full, 2 minutes after the completion of the backup/archive operation. If any jobs are initiated within the 2 minutes, they are allowed to write to the media. If the job was associated with the prior media, new media (such as a new tape) will be used for subsequent jobs. (Applies to all backup/archive types.)

Allow other Schedule to use Media Set

This option allows jobs that are part of a schedule policy or schedule and using a specific storage policy to start a new media and also prevent other jobs from writing to the set of media. It is available only when the **Start New Media** and **Mark Media Full** options are enabled, and can be used in the following situations:

- When one scheduled job initiates several jobs and you only want to start new media on the first job.
- When you want to target specific backups to a media, or a set of media if multiple streams are used.

Extend Job Retention

- **Infinite:** Select this option to retain this job indefinitely.
- **Number of Days:** Select this option to prune this job after the number of days specified.
- **Storage Policy Default:** Select this option to apply the retention rules of the associated storage policy, which is the default option.

Advanced Backup Options (Data Path)

Select the data path to be used for the backup/archive operation.

Ensure that the Library, MediaAgent, Drive Pool, and the Drive selected for this operation is available online and is a part of the associated Storage Policy.

Use Library

Specifies the name of the library that will be used to perform the backup operation. Use this option when you wish to backup to a specific library.

Use MediaAgent

Specifies the name of the MediaAgent that will be used to perform the backup operation. If necessary, you can change the name of the MediaAgent.

For example, if the library is shared and you wish to use a specific MediaAgent (instead of the system selected MediaAgent, or a MediaAgent which may be idle, or less critical) or if you know that the library attached to the specified MediaAgent.

Use Drive Pool

Specifies the name of the Drive Pool that will be used to perform the backup operation. Use this option when you wish to backup using a specific Drive Pool.

Use Drive

Specifies the name of the Drive that will be used to perform the backup operation. Use this option when you wish to backup using a specific Drive from the selected Drive Pool.

Vault Tracking

Select options to export and track media, using Vault Tracker.



Vault Tracking Options will be displayed only when a Vault Tracker license is available in the CommServe.

Vault Tracking options are only applicable for data protection operations using a storage policy associated with a library containing removable media (e.g., tape, optical or stand-alone.)

Export media after the job finishes

Specifies the media used by the data protection operation and media with the specific Media Status (if specified) will be exported and tracked by Vault Tracker.

Exclude Media Not Copied

When selected, allows you to exclude media with jobs that have to be copied.

Media Status

- **All**
Click to select all media. Clear this option to select media with a specific status.
- **Active**
Click to select media with its status marked as active.
- **Full**
Click to select media with its status marked as full.
- **Overwrite Protected**
Click to select media with its status marked as read-only .
- **Bad**
Click to select media with its status marked as bad.

Export Location

Specifies the destination location and lists the stationary locations entered using the **Export Location Details** dialog box.

Track Transit

Specifies that transit information must be tracked, and lists the transit locations entered using the **Export Location Details** dialog box.

Use Virtual Mail Slots

Specifies the exported media is stored within the library in the virtual mail slots defined in the **Library Properties (Media)** dialog box.

Filter Media By Retention

Specifies that the system must automatically filter media based on whether the media has extended retention jobs or not.

- **Media with extended retention job(s)**
Specifies that media with at least one extended retention job will be exported.
- **Media with no extended retention job(s)**
Specifies that media with no extended retention jobs will be exported.

Alert

Use this tab to configure an alert for a schedule policy.

Configure Alert

- **Alert**

The currently configured Alert.

- **Add/Modify Alert**

When clicked, opens the Alert Wizard to configure alerts for this operation.

- **Delete Alert**

When clicked, deletes any existing alerts that are already configured.

SQL Restore Options (General)

Use this dialog box to control how the specified data is to be restored.

Note that restore options vary for single and multiple database restore jobs. The differences are noted below:

Destination Server

Displays the name of the instance to which the selected data will be restored. To change the destination, click one from the list.

The list includes SQL Servers in the CommCell which have the SQL Server iDataAgent installed. Note the following:

- SQL Server 2000 databases can be restored to servers with SQL Server 2000 or SQL Server 2005.
- SQL Server 2005 databases can be restored to servers with SQL Server 2005.

Source Server

Displays the name of the source server or database, depending on whether an instance or database was selected for restore.

Database List

Displays the names of the databases for restore if there is more than one such database.

Server State (or Database State)

Displays the current state of the server or database, depending on whether an instance or database was selected for restore.

Restore Types

Specifies the type of restore job.

• Database Restore

Specifies that the restore operation must restore the selected databases beginning with the most recent full backup and then applying differential and/or transaction log backups up to the selected restore time.

• Step Restore

Specifies that the restore operation must perform a step restore by applying the selected transaction log.

Note that when you have a Quick Recovery Agent, transaction logs for a QR Volume that contains SQL data can also be restored using this option.

Whether the logs are applied to a recovered Quick Recovery Agent volume or to a SQL step restore operation, this option is available only under the following conditions:

- Single database restores
- The restore destination is the source computer
- The database backup and restore histories have not been removed from the SQL server's history tables
- The selected database was previously restored to:
 - STANDBY state for SQL step restore
 - NORECOVER state for QR Agent log application

This option is not available for:

- Instance restores
- Multiple database restores
- File/File Group restores
- When restoring data protection operations that used VSS to create the snapshot

• Recover Only

Specifies that the restore operation must recover the databases to an available online state. No data is actually written. All other options are disabled when this option is selected.

• Point in time

Specifies that the restore operation must restore single or multiple databases to the selected date and time.

Note the following when this option is selected:

- In order to restore a database to a point in time, a transaction log backup must exist.
- If a single database is restored, then the **Restore Time** selection changes from a list of available backups to selection boxes for date and time
- If multiple databases are restored, then the **Restore Time** selection is already in the date/time selection format and clicking **Point-in-Time** accepts the selected date and time exactly as selected. If the restore time falls within the time span of a transaction log backup, the database is restored to the exact time indicated. Otherwise, the database is restored to the latest available backup prior to the selected time.

If the selected point in time is before any database backup was performed, a time selection out of range message will be displayed.

- **Transaction Log Marks**

Specifies that the restore operation must restore transaction marks in the transaction log. This option is enabled only when the database(s) selected for restore have transaction marks in the transaction log. The **Restore Time** list displays a list of transaction marks that exist in the selected database(s) for the given restore time. When multiple databases are being restored, this list shows only those marks that are common in name and time stamp to all selected databases in order to effect an in-sync restore.

- **Partial Restore**

Partial restores, known as "Piecemeal Restore" in SQL Server 2005, allow databases that contain multiple filegroups to be restored and recovered in stages. The process starts with a restore of the primary filegroup, then a restore or recovery of one or more secondary filegroups.

- **Latest Backup Data**

Allows restores of the latest backup data on one or more **non system** secured database in the SQL Server. You can use this feature to create a spare copy of the primary SQL server database (hot standby server) within the same domain, a different domain, or across a Wide Area Network (WAN).

- **Apply Log backups only**

Only available when the **Latest Backup Data** option is selected. Select this option if you wish to include only the transaction log backups completed since the last restore operation. The restore operation will then skip any full backup jobs completed since the last restore operation.



If you do NOT select the **Apply Log Backups Only** option and the database already exists, all backup jobs completed since the last restore operation will automatically be restored. If you do NOT select the **Apply Log Backups Only** option and the database does not already exist, the software will automatically perform a full restore of the database.

- **Skip Full Backup**

Select this option to NOT restore the last full backup since the specified date and time, but to instead restore only the transaction log backups. This option can be useful if the latest full backup has already been restored by third-party software and thus a restore of the SQL iDataAgent full backup is not necessary. Note that in such a case, this option only needs to be selected if a SQL iDataAgent full backup exists which is later than the full backup currently restored on the client by a third-party software.

Restore Time and Time Zone

Specifies the restore time depending on the selected options:

- For a multiple database or a full instance restore, this option indicates the time to which you want to restore all selected databases.
- For a single database restore, this option appears as a list of available backups from which you can select to restore.
- For a SQL Step Restore, this option appears as a list of available transaction log backups from which you can select a single log for restore.

For a QR Volume Step Restore, this option appears as a list of available transaction log backups from which you select a log for restore.

While the SQL Server iDataAgent functionality applies one log only, for the QR Volume the step restore applies the selected log and all logs that occurred prior to the selected log.

- **Restore Time**

Provides a space for you to enter the date and time. Click the arrow to display a calendar.

- **Time Zone**

Lists the time zones. To change the time zone, select one from the list.

- **Stop at mark**

Specifies that the restore operation must restore to the point of the marked transaction. This option is displayed when Transaction Marks are restored.

- **Stop before mark**

Specifies that the restore operation must restore to the point when the marked transaction is committed. This option is displayed when Transaction Marks are restored.

Recovery Type and Undo Path

Specifies the type of recovery that must be performed. This is the state of the database after the restore.

- **Recovery**

Specifies that the database will be fully operational.

- **Norecovery**

Specifies that the database will be offline.

- **Stand By**

Specifies that the database will be in the read-only state (transaction logs may be applied)

- **Undo File Path**

Displays the path in which the undo file is stored, when the **Recovery Type** is **Stand By**. Use the space to modify the default path.

Once the restore is complete, the SQL Server iDataAgent will use the data from the undo file and the transaction log to continue restoring the incomplete transactions. Hence, once the restore completes, the undo file will be re-written with any transactions that are incomplete at that point.

- **Browse** - Click to select a path for the Undo File Path.

Restore Options

- **Unconditionally overwrite existing databases or files**

When this option is checked, data being restored is unconditionally written to the specified location and will overwrite the files of any database that is currently using the database name specified by you. This implements the T-SQL REPLACE command.

When the option is unchecked, restore prevents the accidental data files overwrite by failing the job.

- **Leave database in DBO use only state**

Specifies whether the restore operation leaves database in the DBO_ONLY state where only the owner of the database can access the database.

- **Preserve Replication Settings**

Specifies to restore backups of replicated databases. A published database can be restored to a server other than the server where the database was created. This is relevant only to full backups and if the database database was part of a replication implementation. This option is available only when the **Recovery** type option is selected.

- **Keep Change Data Capture**

Specifies to restore SQL 2008 databases if they were backed up with the Change Data Capture property enabled. This property provides the ability to query a database for changed data. This is a useful SQL feature in that all change activity (insert, update, and delete) is captured and applied to SQL tables. The details of the changes are available in an easily consumed relational format. This option does not apply to backups that were VSS-enabled or to SnapProtect backups.

Advanced

Click to select additional restore options.

Save As Script

Click to open the Save As Script dialog, which allows you to save this operation and the selected options as a script file (in XML format). The script can later be executed from the Command Line Interface using `qoperation execute` command.

When you save an operation as a script, each option in the dialog will have a corresponding xml parameter in the script file. When executing the script, you can modify the value for any of these XML parameters as per need.

To view the XML values for each of the options in the dialog, see the following:

Command Line XML Options

[Back to Top](#)

Job Initiation

Select from the following options. Note that all the options described in this help may not be available and only the options displayed in the dialog box are applicable to the operation for which the information is being displayed.

Immediate

Run This Job Now

Specifies this job will run immediately.

Schedule

Specifies this job will be scheduled. When you click **Configure**, the Schedule Details dialog box will open and allow you to configure the schedule pattern.

Result file location on CommServe

This option is only applicable for List Media operations. Specifies the name of the file in which the results of the scheduled list media job must be saved. Click **Browse** to access to the **Directory Browse** dialog box which allows you to select the folder / file in the CommServe computer.

Automatic Copy

Specifies that an auxiliary copy operation will be performed at the interval specified. This operation will occur when new data that must be copied is found on the primary copy.

Interval

The time interval in which the Automatic Copy will be performed. The default interval is 30 minutes.

Save As Script

Click to open the Save As Script dialog, which allows you to save this operation and the selected options as a script file (in XML format). The script can later be executed from the Command Line Interface using `qoperation execute` command.

When you save an operation as a script, each option in the dialog will have a corresponding xml parameter in the script file. When executing the script, you can modify the value for any of these XML parameters as per need.

To view the XML values for each of the options in the dialog, see the following:

Operations Supporting Save As Script

Configure Alert

Provides the necessary options to configure the alerts associated with this operation.

- **Add/Modify Alert**

When clicked, opens the Alert Wizard to configure necessary alerts for this operation.

- **Delete Alert**

When clicked, deletes any existing alerts that are already configured.

Advanced

Click to select additional options.

Script Preview

Click to display the restore script, based on the selected restore options, that will be submitted to RMAN when restores are performed for the selected Oracle client.

Advanced Restore/Recover/Retrieve Options (General)

Use this dialog box to access additional restore/recover/retrieve options. Note that all the options described in this help may not be available and only the options displayed in the dialog box are applicable to the component installed on the client.

Use hardware revert capability if available

Specifies whether to revert the data to the time when the snapshot was created. Selecting this option brings back the entire LUN to the point when the snapshot was created, overwriting all the modifications to data since the snapshot creation. This option is only available if the snapshot engine used for SnapProtect backup supports the revert operation.

Use RMAN Restore

Specifies whether to use RMAN for the restore operations.

Use FileSystem Restore

Specifies whether to use file system for the restore operations.

Advanced Restore Options (Pre/Post)

Establish pre/post processes for restore jobs, and the account that has permissions to run these processes (for Windows-based agents).

Pre Recovery Command:

Displays the name/path of the process to run before the restore. Add or modify the name/path, or use the **Browse** button to search for and select the name/path. The system allows the use of spaces in the name/path, provided they begin with an opening quotation mark and end with a closing quotation mark.

Post Recovery Command:

Displays the name/path of the process to run after the restore. Add or modify the name/path, or use the **Browse** button to search for and select the name/path. The system allows the use of spaces in the name/path, provided they begin with an opening quotation mark and end with a closing quotation mark.

Run Post Restore Process for all attempts

Specifies whether this process will execute for all attempts to run the phase. Selecting this option will execute the specified process for all attempts to run the phase, including situations where the job phase is interrupted, suspended, or fails. Otherwise, when the checkbox is cleared the specified process will only execute for successful, killed, or failed jobs.

Pre/Post Impersonation

For Windows-based agents, you must designate either the Local System Account or, for added security, another account as having permission to run these commands for restore jobs.

- **Use Local System Account**

Normally, the Local System Account has permissions to access all the data on the local computer.

- **Impersonate User**

Select this check box to enable the User Name and Password boxes. If the Impersonate User account defined here is not available, restore jobs using pre/post commands will fail. This account operates independently of the Impersonate User account for backup jobs.

- **User Name**

Enter the Window's user account name which will have permission to execute the desired commands.

- **Password**

Enter the corresponding password for this account.

- **Confirm Password**

Enter the password again for this account.

Advanced Restore Options (Copy Precedence)

Choose the copy from which you wish to recover or retrieve. Select from the following options:

Restore from copy precedence

When selected, the system retrieves the data from the storage policy copy with the specified copy precedence number. If data does not exist in the specified copy, the operation fails even if the data exists in another copy of the same storage policy.

When cleared, (or by default) the system retrieves data from the storage policy copy with the lowest copy precedence. If the data was pruned from the primary copy, the system automatically retrieves the data from the other copies of the storage policy starting with the copy with the lowest copy precedence and proceeding through the copies with higher copy precedence. Once the data is found, it is retrieved, and no further copies are checked.

Copy Precedence

Specifies the precedence number with which the system recovers or retrieves data from the copy.

Advanced Restore Options (Data Path)

Select the data path for the restore/recovery operation. You can specify the MediaAgent, Library, Drive Pool, and Drive from which the restore operation must be performed.

Use MediaAgent

Specifies the name of the MediaAgent that will be used to perform the restore operation. If necessary, you can change the name of the MediaAgent.

For example, if the library is shared and you wish to use a specific MediaAgent (instead of the system selected MediaAgent, or a MediaAgent which may be idle, or less critical) or if you know that the media containing the data you wish to restore is available in the library attached to the specified MediaAgent.

If the media containing the data is not available in the tape/optical library attached to the MediaAgent, the system will automatically prompt you to insert the appropriate media. In the case of a disk library, the operation will fail if the requested data is not available in the disk library attached to the specified MediaAgent.

Use Library

Specifies the name of the library that will be used to perform the restore operation. Use this option when you wish to restore using a specific library.

For example, if you know that the media containing the data you wish to restore is available in a specific library.

Use Drive Pool

Specifies the name of the Drive Pool that will be used to perform the restore operation. Use this option when you wish to restore using a specific Drive Pool.

To restore NAS data, select the drive pool type that was used to perform the backup, *i.e.*, if a drive pool associated with an NDMP Remote Server was used to perform the backup, select a drive pool associated with an NDMP Remote Server. Similarly, if an NDMP drive pool was used, specify an NDMP drive pool.

Use Drive

Specifies the name of the drive in the drive pool that will be used to perform the restore operation. Use this option when you wish to restore using a specific Drive in the Drive Pool.

Use Proxy

Specifies the name of the proxy server that will be used to perform the restore operation. Use this option when you wish to restore using a proxy server.

Advanced Restore Options (Encryption)

Pass-Phrase

Enter the pass-phrase that is currently assigned to the client, whose data you are restoring. Note that if you have changed the pass-phrase since you secured the client data, you need to provide the new pass-phrase here, not the old one.

Re-enter Pass-Phrase

Re-enter the pass-phrase for confirmation.

If you attempt an immediate restore of encrypted data that was pass-phrase protected without entering the pass-phrase here, the restore operation will fail.

If you have an exported pass-phrase set up, and you enter the pass-phrase under **Decryption**, you over-ride (not overwrite) the client properties pass-phrase. Thus, if you enter the pass-phrase incorrectly, the restore does not complete successfully.

Advanced Restore Options (Options)

Use this dialog box to make modifications in the names and locations of databases for restore.

Database

Displays the name of the database. You can edit the name by clicking within the box. Changes cause the data to be restored to a new database on the selected server.

For a detailed description on how to change the name of a SQL database, see the procedure titled "Restore a Database with a Different Name" in *Books Online*.

Device Name

Displays the device name of the database.

File Name

Displays the name of the database file to be restored.

Physical Path

Displays the complete path to the selected database file. You can edit the path by clicking within the box. Changes cause the database files to be restored to a new location. If the specified path does not exist, it will be created during the restore process.

For SQL, if the database file path is changed without changing the name of the database, then the existing database will be overwritten and will point to the new location. For a detailed description, see the procedure titled "Restore with a Different Data File Path" in *Books Online*.

Size (MB)

Displays the size of the database file in megabytes.

Find & Replace

For SQL, click to change multiple Database Name or Database File Path.