

Administration Guide - Documentum *iDataAgent* (DB2)

TABLE OF CONTENTS

OVERVIEW

Introduction

Key Features

- Supported Environments
- Full Range of Backup and Recovery Options
- Scheduling
- Efficient Job Management and Reporting
- Restore Levels

Terminology

SYSTEM REQUIREMENTS - DOCUMENTUM /DATAAGENT

SUPPORTED FEATURES - DOCUMENTUM /DATAAGENT

DEPLOYMENT - DOCUMENTUM /DATAAGENT

INSTALL THE DOCUMENTUM /DATAAGENT - UNIX

INSTALL 32-BIT DOCUMENTUM /DATAAGENT ON A MICROSOFT WINDOWS X64 PLATFORM

INSTALL 32-BIT DOCUMENTUM /DATAAGENT - WINDOWS

INSTALL 64-BIT DOCUMENTUM /DATAAGENT - WINDOWS

CONFIGURATION - DOCUMENTUM /DATAAGENT - BASIC

Overview

Create an Instance

Configure a Subclient

CONFIGURATION - DOCUMENTUM IDATAAGENT - ADVANCED

Overview

Planning Your Backups

- Storage Area Backup
- Archive Log Backup
- Offline Database Backup
- Full-Text Index Backup
- Online Database Backup
- Selective Full Backup

Modifying an Agent, Instance, or Subclient

Deleting an Instance or Subclient

- Delete an Instance
- Delete a Subclient

BACKUP - DOCUMENTUM /DATAAGENT - BASIC

Overview

What Gets Backed Up

Your first Full Backup

- Running a Full Backup Immediately
- Scheduling a Full Backup

Viewing Backup Job History

BACKUP - DOCUMENTUM /DATAAGENT - ADVANCED

Overview

Incremental backups

Starting an Incremental Backup

Database Backups

Offline Backups

Online Backups

Archive Log Backups

Selective Copy - Full Backups

Enable Database Configuration Parameters

Storage Area Backups

Full-text Index Backups

Managing Jobs

Restarting Jobs

Controlling Jobs

Additional Options

RESTORE - DOCUMENTUM /DATAAGENT - BASIC

Overview

Your First Restore

Documentum File Browse and Restore

Component-Level Browse and Restore

Viewing Restore Job History

Viewing Restored Items Associated With a Restore Job

RESTORE - DOCUMENTUM /DATAAGENT - ADVANCED

Overview

Component-Level Browse and Restore Operations

Restoring Storage Areas

Restoring Full-Text Indexes

Restoring the Entire Database

Restoring the Database while Online/Offline

Restoring History Files

Restoring Logs

Reapplying Log Files with Roll-Forward

Rename one or more Tablespaces with Redirect

Recovering a Database

Documentum File Browse Operations

Restoring a Single Document

Restoring Multi-Versioned Documents

Restoring a Document to a Different Cabinet

Command Line Restores

Browse Data

Selecting Objects From the Browse Window for Restore

Controlling the Browse Time Interval

Browsing Data Associated with Specific Jobs

Global Browse / Restore Options

Browse Error Reporting

Browsing and Restoring Using a Specific MediaAgent, Library or Drive Pool

Browsing and Restoring From Copies

List Media

Listing Media For a Subclient

Listing Media For a Backup Set or Instance

Restore by Jobs

Manage Restore Jobs

Restart Jobs

Resubmit Jobs

Control Jobs

Additional Options

DISASTER RECOVERY - DOCUMENTUM /DATAAGENT

Overview

Planning for a Disaster Recovery

Oracle Database Considerations

View Application Version

Perform a Full System Restore

PHASE I. Restore Documentum Components

PHASE II. Restore Database Components

PHASE III. Restore Storage Areas, Full-Text Index Components

DATA AGING

Overview

Configuring Retention Rules

Basic Retention Rules

Selective Online Full Backup Retention Rules

Documentum Command Line Retention Rules

Archive Log Retention Rules

Extended Retention Rules

Running or Scheduling Data Aging Jobs

Viewing Data Aging Job History

ADDITIONAL OPERATIONS

Overview

Related Reports

Creating a Report

Audit Trail

Setting up Audit Trail

Policies

Storage Policy

Subclient Policy

Schedule Policy

Auxiliary Copy

Starting an Auxiliary Copy

Operation Window

Defining Operation Rules

Operating System and Application Upgrades

Online Help Links

Uninstalling Components

General Uninstall Considerations

Uninstalling Unix Components

Uninstall Software Using the CommCell Console

FREQUENTLY ASKED QUESTIONS

TROUBLESHOOTING

Backup

Changes to Full-Text Index Path in Instance Properties is not Reflected in Subclient Restore

Restore

Troubleshooting Oracle Errors

Overview - Documentum iDataAgent

TABLE OF CONTENTS

Introduction

Key Features

- Supported Environments
- Full Range of Backup and Recovery Options
- Scheduling
- Efficient_Job_Management_and_Reporting
- Restore Levels

Terminology

INTRODUCTION

The Documentum iDataAgent provides a data protection and recovery solution to protect and manage EMC Documentum® Repository data, such as databases, Storage Areas, and Full-Text Indexes (FTI). Documentum Repository with Oracle and DB2 databases are supported. The database components that are backed up include data files and logs, which hold the Repository's meta data.

Use the Documentum iDataAgent for the following reasons:

- Backup of databases, Storage Areas, and Full-Text index in a Single or Distributed Server environment.
- Restore of databases, Storage Areas, and Full-Text index in a Single or Distributed Server environment.
- Restore individual files or documents of a Repository to an in-place or out-of-place destination within the same Repository.

KEY FEATURES

The Documentum iDataAgent offers the following key features:

SUPPORTED ENVIRONMENTS

As the following sections discuss, both single-server and distributed-server environments are supported by the Documentum iDataAgent.

SINGLE-SERVER ENVIRONMENT

A single Repository is supported in a single-server, all-in-one configuration. When all Documentum components are installed on the same computer, the Documentum iDataAgent is sufficient to backup all components.

The illustration to the right depicts a possible example of how you could configure your environment. Keep in mind that configuration is entirely flexible and that you should set up your client computers and install your components according to your needs and available resources.

In the single-server configuration, the Documentum iDataAgent is installed on the same computer as the Primary Documentum Content Server to protect these components:

- Storage Areas
- Database (Oracle or DB2)
- Full-Text Indexes (optional)

Contrast a single-server environment with a Distributed-Server Environment.

DISTRIBUTED-SERVER ENVIRONMENT

In addition to a single-server environment, a single Repository is supported in a distributed-server configuration. When the components are installed in a distributed setup, the File System iDataAgent or Database iDataAgents (Oracle or DB2) need be installed according to the configuration.

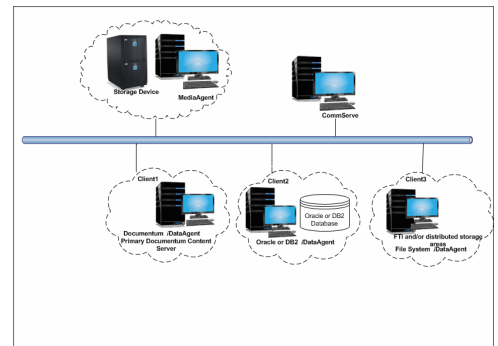
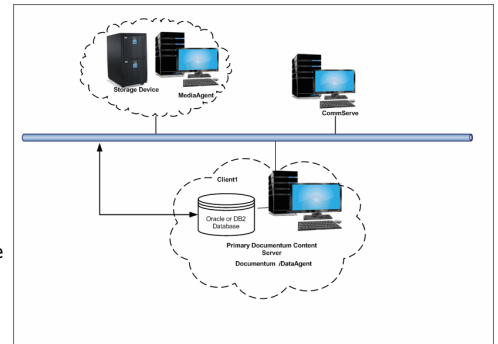
The illustration to the right depicts a possible example of how you could configure your environment. Keep in mind that configuration is entirely flexible and that you should set up your client computers and install your components according to your needs and available resources.

In this configuration, these components are installed on three client computers as follows:

• Client 1

This is the Primary Content Server where Documentum iDataAgent must be installed to protect the Storage Area residing on the server.

• Client 2



On this computer, a DB2 or Oracle iDataAgent is installed to protect the following the database components (Oracle or DB2).

- **Client 3**

On this computer, the File System iDataAgent is installed to protect the following components:

- Distributed Storage Areas
- Full-Text Index

Contrast this with a Single-Server Environment where all Documentum components are installed on the same computer and the Documentum iDataAgent is sufficient to backup all components.

FULL RANGE OF BACKUP AND RECOVERY OPTIONS

Documentum iDataAgent provides the flexibility to backup Storage Areas, files, and database (Oracle or DB2) components in an online, zero-downtime backup environment. This is very essential since database data is very dynamic and is subject to constant changes.

You can perform a full backup for Storage Areas, Full-Text Indexes, and database (Oracle or DB2) components. You can also perform an incremental backup of database and Storage Areas. The following section describes the backups that can be performed in different environments.

STORAGE AREA BACKUP

The Documentum iDataAgent enables you to back up both local and distributed Storage Areas that are distributed across different computers. With both distributed and local storage areas, the File System iDataAgent should be installed on the client computers that are hosting the Storage Areas. Both full and incremental backup types are supported for Storage Areas. The first backup of a Storage Area is performed as a Full backup. Storage Areas can be backed up in a full Repository backup or as a separate backup.

FULL-TEXT INDEX BACKUPS

The Documentum iDataAgent enables you to back up Full-text Indexes. The File System iDataAgent should be installed on the client computers that are hosting the Full-text Indexes. Full-text Indexes can be backed up only in full backups. Full-text Indexes can be backed up in a full Repository backup or as a separate backup.

When backing up a Full-Text Index, it necessary to quiesce the Index Server to ensure consistent backups. For this purpose, quiesce and unquiesce scripts are required. These will quiesce and unquiesce the Indexer process in the Index Server during the backup operation.

These two sample scripts are provided in the `<installation_folder>\Base` folder on the computer where the Documentum iDataAgent has been installed:

- `cvfti_startup.sh`
- `cvfti_shutdown.sh`

These scripts have to be modified to suit the environment prior to back up.

DATABASE OFFLINE BACKUP

For both DB2 and Oracle databases, when the database is shutdown and not available for use, you can perform a full backup of the database without the logs. This is especially used when the data is consistent and there are no transactions in the database.

DATABASE ONLINE BACKUP

For both DB2 and Oracle databases, in cases, when you cannot bring down the database to perform an offline backup, you can use the online backup method. Here, you can perform full or incremental backups when the database is online. For Oracle databases, it must be in ARCHIVELOG mode, especially when you want to perform a point-in-time recovery of the Oracle Database.

You also have the facility the backup only the archive logs when the database is online. These logs can be applied to an online backup to recover the database to the current point-in-time.

You can also protect the non-database files and profiles using the File System iDataAgent.

SCHEDULING

Schedules can be set up to automate data protection and data recovery operations. Scheduling jobs helps to ensure that operations required to manage the CommCell are automatically performed on a regular basis without user intervention, which is extremely beneficial in managing a CommCell.

EFFICIENT JOB MANAGEMENT AND REPORTING

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

RESTORE LEVELS

You can turn back the clock to a specific point-in-time and retrieve desired data. Restores can be run on-demand or can be scheduled to be run at a later point of time. When restoring data, the Documentum iDataAgent provides the facility to restore the data to the same path/destination on the client computer from which the data originated; this is referred to as an in-place restore. All restore levels provide the facility to restore data to a different directory on the client computer from which the data originated; this is referred to as an out-of-place restore to a different path/destination.

TERMINOLOGY

The Documentum iDataAgent documentation uses the following terminology:

FULL-TEXT INDEX	A Full-Text Index is an optional component. The Content Server maintains a Full-Text Index for all content stored in the repository. The Full-Text Index enables rapid search through a large collection of unstructured information. It searches for content and attributes in a given index and can be setup to run on multiple hosts.
OFFLINE BACKUP	A backup of database objects performed while the corresponding database instance is shutdown and unavailable to users.
ONLINE BACKUP	A backup of database objects performed while the corresponding database instance is running and available to users.
QUIESCE	Prior to a backup, a quiesce operation must occur for a Documentum Full-Text Index.
REPOSITORY	A Documentum term that is often synonymous with <i>Docbase</i> . It is made up of the following components: <ul style="list-style-type: none"> • One or more Storage Areas for content files • A third-party database for metadata • An optional Full-Text Index These components may reside all on the same server, or on two or more servers.
STORAGE AREA	These hold one or more file stores under the operating system file system. These are the actual content files that are stored in the Storage Area. A Storage Area is typically a file system subdirectory hierarchy of possibly millions of files in all formats. The Documentum iDataAgent currently supports backing up and restoring the File Store Storage Area type. Other types of Storage Areas (e.g., BLOB stores, Linked stores) are not supported.
UNQUIESCE	After a backup, an unquiesce operation occurs for a Documentum Full-Text Index.

[Back to Top](#)

System Requirements - Documentum iDataAgent

System Requirements | Supported Features

The following requirements are for the Documentum iDataAgent:

APPLICATION/OPERATING SYSTEM		PROCESSOR
DOCUMENTUM 6.5 SP1/HIGHER ON:		
AIX	AIX 6.1	Power PC (Includes IBM System p)
HP-UX	HP-UX 11i v3 (11.31)	Itanium
LINUX	RED HAT ENTERPRISE LINUX/CENTOS	
	Red Hat Enterprise Linux/CentOS 5.2 with glibc 2.5.x	Itanium, x64 or compatible processors
	SUSE LINUX (SLES)	
	SuSE Linux 11.x with glibc 2.9.x and above	Itanium, x64 or compatible processors
	SuSE Linux 10 SP2 Enterprise Server	Itanium, x64 or compatible processors
SOLARIS	Solaris 10.x with a minimum of SunOS (Sparc) Patch 119963-14	Sparc5 or higher recommended
WINDOWS SERVER 2003	Microsoft Windows Server 2003 32-bit and x64 Editions with a minimum of Service Pack 2	All Windows-compatible processors supported
WINDOWS	WINDOWS 2008	
	Microsoft Windows Server 2008 32-bit and x64 Editions*	All Windows-compatible processors supported
	*Core Editions not supported	

HARD DRIVE

230 MB minimum of hard disk space for software

MEMORY

64 MB RAM per stream/drive minimum required beyond the requirements of the operating system and running applications

Swap space = 2*RAM size

SUPPORTED DATABASES

ORACLE

Oracle 10g R2 (10.2.x)

Oracle 11g R1

Oracle 11g R2 (supported since Documentum 6.7)

See System Requirements - Oracle iDataAgent for more information on the supported versions listed above.

DB2

DB2 9.5 Universal Database

See System Requirements - DB2 iDataAgent for more information on the supported version listed above.

SQL

MS SQL 2005 SP3

See System Requirements - Microsoft SQL Server iDataAgent for more information on the supported version listed above.

SOLARIS ZONES/CONTAINERS SUPPORT

Data Protection of data residing on global and non-global zones is supported.

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

AIX LPAR/WPAR SUPPORT

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

PERIPHERALS

DVD-ROM drive

Network Interface Card

MISCELLANEOUS

NETWORK

TCP/IP Services configured on the computer.

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

- System Requirements - AIX File System iDataAgent
- System Requirements - HP-UX File System iDataAgent
- System Requirements - Linux File System iDataAgent
- System Requirements - Solaris File System iDataAgent

SELINUX

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

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

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

File Name	Content of the File
<p><directory>/<file_name>.te</p> <p>where:</p> <p><directory> is /usr/share/selinux/devel</p> <p><file_name> is the name of the Unix file, created to save the policy module statement. It is a good idea to use the same name for policy module and the file.</p> <p>For example: When you are creating a policy module for backup_IDA application, you can use the following file name: backup_IDA.te</p>	<p>The content of the file should be as follows:</p> <pre>policy_module(<name>,<version>) ##### where: <name> is the name of the policy module. You can give any unique name to the policy module, such as a process or application name. <version> is the version of the policy module. It can be any number, such as 1.0.0. For Example: While creating a policy module for the backup_IDA application, you can use the following content. policy_module(backup_IDA,1.0.0)</pre>
<p><directory>/<file_name>.fc</p> <p>where:</p> <p><directory> is /usr/share/selinux/devel</p> <p><file_name> is the name of the Unix file, created to save the policy module statement. It is a good idea to use the same name for policy module and the file.</p> <p>For example: When you are creating a policy module for backup_IDA application, you can use the following file name: backup_IDA.fc</p>	<p>The content of the file should be as follows:</p> <p>Note that the following list of files is not exhaustive. If the process fails to launch, check /var/log/messages. Also, if required, add it to the following list of files.</p> <pre>/opt/<software installation directory>/Base/libCTreeWrapper.so -- gen_context (system_u:object_r:texrel_shlib_t,s0) /opt/<software installation directory>/Base/libCVMAGuiImplgso -- gen_context (system_u:object_r:texrel_shlib_t,s0) /opt/<software installation directory>/Base/libdb2locale.so.1 -- gen_context (system_u:object_r:texrel_shlib_t,s0) /opt/<software installation directory>/Base/libdb2osse.so.1 -- gen_context (system_u:object_r:texrel_shlib_t,s0) /opt/<software installation directory>/Base/libDb2Sbt.so -- gen_context (system_u:object_r:texrel_shlib_t,s0) /opt/<software installation directory>/Base/libdb2trcapi.so.1 -- gen_context (system_u:object_r:texrel_shlib_t,s0) /opt/<software installation directory>/Base/libDrDatabase.so -- gen_context (system_u:object_r:texrel_shlib_t,s0) /opt/<software installation directory>/Base/libIndexing.so -- gen_context (system_u:object_r:texrel_shlib_t,s0) /opt/<software installation directory>/Base/libSnooper.so -- gen_context (system_u:object_r:texrel_shlib_t,s0)</pre>

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

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

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

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

SELinux is now configured to work with this application.

DISCLAIMER

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

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

Configuration - Documentum iDataAgent (DB2)

Basic | **Advanced**

TABLE OF CONTENTS

Overview

Configure the DB2 Database Backup

Create an Instance

Configure a Subclient

OVERVIEW

The Documentum iDataAgent uses two components to facilitate backups:

- An **Instance**, which defines the Documentum Repository to be backed up. As soon as you install the agent, you need to manually create an instance using the procedure outlined in Create an Instance before you run the first backup operation.
- A **Default Subclient**, which is automatically created after configuring an Instance and contains all the objects of the instance, such as:
 - DB2 database files, history files, and log files
 - Storage Areas
 - Full-text Indexes (optional)

Before running backups, ensure that the Storage Policy is configured on the subclient and for Full-Text Indexes, if available, Quiesce/UnQuiesce scripts are provided. Use the procedure outlined in Configure a Subclient to edit the contents of the default subclient or create a new one. See the Advanced configuration for a list of possible reasons to create subclients.

CONFIGURE THE DB2 DATABASE BACKUP

Prior to configuring the Documentum iDataAgent, perform the following:

SINGLE SERVER SETUP

If the DB2 database is installed on a separate client from that of the Documentum iDataAgent, the DB2 iDataAgent must be installed on that client computer and a DB2 instance must be configured. For more information, refer to Instances - DB2.

DISTRIBUTED SERVER SETUP

If the DB2 database and Documentum Server are installed on the same computer then you must create the following registry keys on the client computer:

REGISTRY KEY	VALUE
sDB2_ARCHIVE_PATH	<software_install_directory>\archive
sDB2_AUDIT_ERROR_PATH	<Software_install_directory>\audit_error
sDB2_RETRIEVE_PATH	<software_install_directory>\retrieve

CREATE AN INSTANCE

This section describes how to create an instance for the Documentum iDataAgent in order to perform backup and recovery operations.

1. From the CommCell Browser, locate the correct client computer node.

Right-click the agent node for which you want the new instance created, click **All Tasks**, and then click **Create New Instance**.
2. Add the following information in the **General** tab:
 - Enter the name (up to 32 characters) for the instance in **Docbase Name**.
 - Click **Configure Password** and specify the Documentum Install Owner User Name and Password in **User Account**. Confirm the Password and click **OK**.
 - Type the Repository installation path in the **DOCUMENTUM Path** box or use **Browse** to locate the path.
 - Type the Repository shared path in the **DOCUMENTUM_SHARED Path** box or use **Browse** to locate the path.
 - Click **Discover**. The Docbase name and User Name and Password will be validated. If successful, the instance is created and the Docbase Version and Docbase Id are automatically populated.

DOCBASE USER ACCOUNT

To perform data protection and recovery operations, the Documentum iDataAgent requires a user account to log on to the related Documentum docbase to access the data. The user account is the Documentum *Install owner* account that was used to install the Documentum software.

The account must already be set up on the client.

3. Click the **Database** tab and add the following information:

- From the **Client** pulldown menu, select the client computer on which the DB2 database resides.
- Type the DB2 instance name in the **Instance** box.
- Type the DB2 database name in the **Database** box.
- Type the installation home path in the **Home** box or use **Browse** to locate the path.
- Click **Configure Password** to specify the DB2 User account used by the system to perform all Data Protection and Recovery operations for this instance.

To perform data protection and recovery operations, the DB2 iDataAgents require a user account with one of the following privileges to access the DB2 application and database:

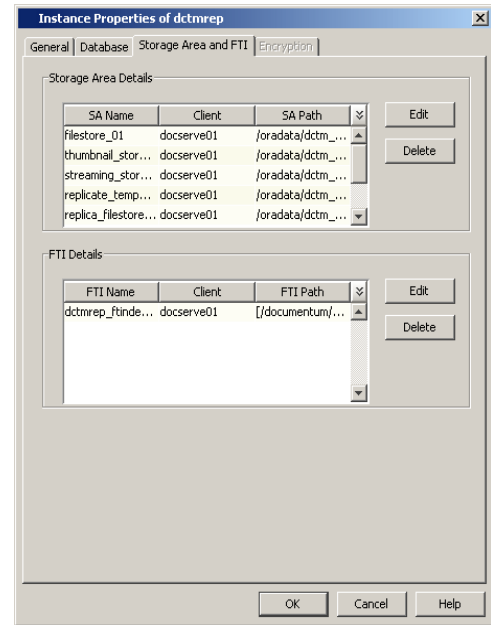
- Administration privileges (default)
- DB2 database administration privileges
- SYSADM, SYSCTRL, and SYMAINT user group privileges

This account must already be set up on the client. Additional accounts should be established by the DB2 database administrator. To establish additional accounts on your own, consult the appropriate DB2 application documentation.

- Click **OK**.

4. Click the **Storage Area and FTI** tab and add the following information:

- In **Storage Area Details**, select a file store that was discovered and click **Edit**. In the Edit Operation dialog box, select the **Client** name from the pulldown menu.
- If a Full-text Index component is present, in **FTI Details**, select the full-text index name that was discovered and click **Edit**. In the Edit Operation dialog box, select the **Client** name from the pulldown menu and type the **FTI Path**.




5. Click **OK** to save the instance.


CONFIGURE A SUBCLIENT

The following procedure describes the steps involved in configuring a subclient. This procedure may be used to configure the default subclient, as well as to create any additional subclients. Refer to Create User-Defined Subclients for more information.

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





2. From the **Full Text Index** tab, you can:
 - Exclude backing up the FTI for this subclient by deselecting **Backup FTI**.
 - Select a **FTI** and click **Edit** to change the client computer or the path in which it resides.
 - If applicable, click **Add** to add a new **FTI**.
 - Change the name or path of the **Quiesce** or **Unquiesce** script.

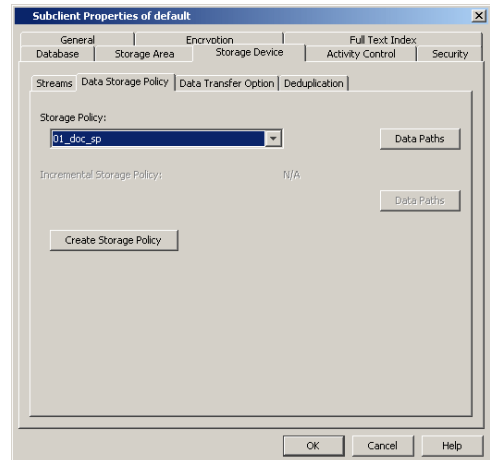


3. From the **Database** tab, you can:
 - Exclude backing up the database for this subclient by deselecting **Backup DB**.
 - Change the state in which the database will be backed up (**Offline** or **Online**).
 - Include or exclude backing up **Archive Logs** or **Deleted Logs**.
 - Specify **Pre** or **Post Backup Process** (if applicable).

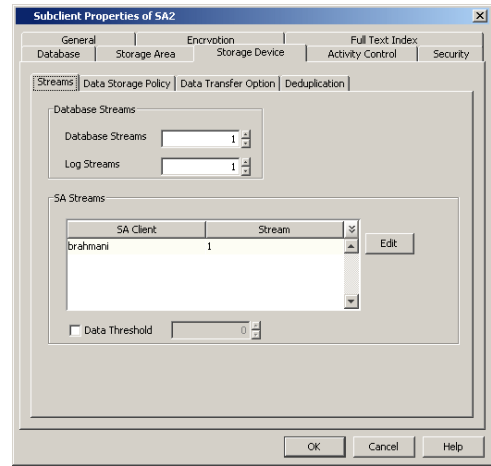
4. From the **Storage Area** tab, you can:
- Exclude backing up the Storage Area for this subclient by deselecting **Backup SA**.
 - Select a **SA** and click **Edit** to change the client computer in which it resides.
 - If applicable, click **Add** to add a new **SA**.
 - Specify a **Pre** or **Post Backup Process** (if applicable).



5. From the **Storage Device** tab, click the **Data Storage Policy** tab and specify a storage policy for the subclient.



6. From the **Storage Device** tab, click the **Streams** tab and specify streams for Database, Logs, and Storage Area.



7. Click **OK** to save the subclient.

[Back to Top](#)

Configuration - Documentum iDataAgent (DB2)

Basic | **Advanced**

TABLE OF CONTENTS

Overview

Planning Your Backups

- Storage Area Backup
- Archive Log Backup
- Offline Database Backup
- Full-Text Index Backup
- Online Database Backup
- Selective Full Backup

Modifying an Agent, Instance, or Subclient

Deleting an Instance or Subclient

- Delete an Instance
- Delete a Subclient

OVERVIEW

Based on your environment, you might need to fine-tune your Documentum instance or subclient. Following are the additional configurations that you can use to configure this Agent.

PLANNING YOUR BACKUPS

In order to recover a database during the event of a failure or data loss, you need to ensure that you always have the latest and consistent copy of the Documentum Repository data. To do this, you need to plan and decide on the backups to be performed. It is always recommended that you create separate subclients for each of these backups.

Although you can re-configure the content of the Default Subclient to backup specific objects, we strongly recommend against it because this would disable the capability of the Default Subclient to serve as a catch-all entity for client data, thus increasing the likelihood that some data will not get backed up.

STORAGE AREA BACKUP

Splitting up storage areas across different subclients allows for complete control of how often they are backed up. For example, you may want to backup specific storage areas four times a day and backup the database at the end of the day.

Therefore, a user-defined subclient must be created to specify only storage areas that are separate from the database.

1. From the CommCell Browser, right-click the instance for which you want to create a new subclient, click **All Tasks** and then click **New Subclient**.
2. Click the **General** tab, type the name (up to 32 characters) of the subclient that you want to create.
3. Click the **Storage Area** tab:
 - o Select the **Backup SA** checkbox.
 - o Click **Add** and enter the **SA Name**, **Client**, and its **SA Path** in which the storage area resides.
4. Click the **Storage Device** tab:
 - o Select a storage policy to associate with this subclient from the storage policy list.
 - o Click the **Streams** tab and configure streams for the storage areas.
5. Click **OK** to save the subclient configuration.
 - Ensure that **Backup FTI** and **Backup DB** are not selected in the **Full Text Index** and **Database** tabs.
 - In order to discover far stores and their associated clients in a distributed environment, ensure that they are configured correctly.

ARCHIVE LOG BACKUP

Specifying only log files in a subclient allows you to back them up up more frequently than backing up the entire database. Archive log backups enable you to recover database transactions that have been lost due to an operating system or disk failure. You can apply these archive logs to an online backup in order to recover a database.

Full backups generally includes both the data and logs. However, because of their importance in recovering data, it is recommended that you create separate subclients for backing up archived log files.

1. From the CommCell Browser, right-click the instance for which you want to create a new subclient, click **All Tasks** and then click **New Subclient**.

2. Click the **General** tab, type the name (up to 32 characters) of the subclient that you want to create.
3. Click the **Database** tab:
 - o Click **Backup Archive Log**.
4. Click the **Storage Device** tab:
 - o Select a storage policy to associate with this subclient from the storage policy list.
 - o Click the **Streams** tab and configure streams for the log.
5. Click **OK** to save the subclient configuration.

Ensure that **Backup FTI** and **Backup SA** are not selected in the **Full Text Index** and **Storage Area** tabs.

OFFLINE DATABASE BACKUP

Offline backups are performed when the database is shut down and unavailable for use. Offline backups are always full backups and do not include the archived log files.

In order to backup the database when it is offline, you need to create a separate user-defined subclient for offline backup.

1. Make sure that the database used for creating the subclient is up and running. This ensures that the system accesses the most recent configuration of the database.
2. From the CommCell Browser, right-click the instance for which you want to create a new subclient, click **All Tasks** and then click **New Subclient**.
3. Click the **General** tab, type the name (up to 32 characters) of the subclient that you want to create.
4. Click the **Database** tab:
 - o Select the **Backup DB** checkbox and click **Offline Database**.
5. Click the **Storage Device** tab:
 - o Select a storage policy to associate with this subclient from the storage policy list.
 - o Click the **Streams** tab and configure streams for the database.
6. Click **OK** to save the subclient configuration.

Ensure that **Backup FTI** and **Backup SA** are not selected in the **Full Text Index** and **Storage Area** tabs.

FULL TEXT INDEX BACKUP

Including only full-text indexes in a subclient allows for complete control of how often they are backed up. For example, you may want to backup full-text indexes periodically.

Therefore, a user-defined subclient must be created to specify only full-text indexes.

1. From the CommCell Browser, right-click the instance for which you want to create a new subclient, click **All Tasks** and then click **New Subclient**.
2. Click the **General** tab, type the name (up to 32 characters) of the subclient that you want to create.
3. Click the **Full Text Index** tab:
 - o Select the **Backup FTI** checkbox.
 - o Click **Add** and enter the **FTI Name**, **Client**, and its **FTI Path** in which the full-text index resides.
4. Click the **Storage Device** tab, select a storage policy to associate with this subclient from the storage policy list.
5. Click **OK** to save the subclient configuration.

Ensure that **Backup DB** and **Backup SA** are not selected in the **Database** and **Storage Area** tabs.

ONLINE DATABASE BACKUP

Online backups are performed when the database is online. Data files are included when you back up the database in online mode. Archived log files can also be included in the backup. These backups do not require database downtime and are extremely useful when you need to perform a point-in-time restore of the DB2 database.

You can perform a full or incremental backups when the database is online. In order to backup the DB2 database when it is online, you need to create a separate user-defined subclient for online backup.

1. Make sure that the database used for creating the subclient is up and running. This ensures that the system accesses the most recent configuration of the database.
2. From the CommCell Browser, right-click the instance for which you want to create a new subclient, click **All Tasks** and then click **New Subclient**.
3. Click the **General** tab and type the name (up to 32 characters) of the subclient that you want to create.
4. Click the **Database** tab:
 - o Select the **Backup DB** checkbox and click **Online Database**.
 - o Click **Backup Archive Log** (optional).
 - o Click **Delete Archive Log** (optional).
 - o Specify a **Post Backup Process** (optional).
5. Click the **Storage Device** tab:
 - o Select a storage policy to associate with this subclient from the storage policy list.
 - o Click the **Streams** tab and configure streams for the database.
6. Click **OK** to save the subclient configuration.

Ensure that **Backup FTI** and **Backup SA** are not selected in the **Full Text Index** and **Storage Area** tabs.

SELECTIVE COPY - FULL BACKUP

A selective copy allows you to copy backup data selectively from a source copy to this copy, providing for better tape rotation. Since only selective backups can be copied to selective copies, the selective copies cannot be promoted to the primary copy, only synchronous copies can be promoted. Note that the data selection process does not have to be the same for all auxiliary copies.

During an auxiliary copy operation or a data protection operation that creates an inline copy, only those backups from the primary copy that meet certain criteria will be copied to a selective copy. You can define a selective copy to be time-based, automatically selected, or automatically not selected on the primary copy.

If the copy is defined as All Fulls, all full backups on the primary copy will be copied during an auxiliary copy operation or a data protection operation that creates an inline copy. If the copy is defined as time-based, only the first or last full backup that occurs within each selected weekly, monthly, quarterly, half-yearly, or yearly interval will be copied. You have the option of creating and associating a custom calendar to the copy, so that the intervals can be further customized.

Selective copy will not only copy Selective Online Full and Selective Offline Full jobs, it will also copy regular Full jobs with the linked Log backup jobs.

CONFIGURING SUBCLIENT FOR SELECTIVE COPY WITH FULL BACKUPS

1. Create a selective copy:
 - o From the CommCell Browser, right-click the storage policy for which you wish to create the secondary copy, click **All Tasks** and then click **Create New Copy**.
 - o From the **General** tab of the **Copy Properties** dialog box:
 - Enter the copy name in the **Copy Name** box.
 - Select the **Selective Copy** check box. The **Selective Copy** tab is enabled.
 - Select the library, MediaAgent, drive pool and scratch pool from the lists (drive pool and scratch pool are not applicable for disk libraries).
2. Associate a subclient with the storage policy copy:
 - o From the CommCell Browser, right-click the instance for which you want to create a new subclient, click **All Tasks** and then click **New Subclient**.
 - o Click the **Subclient Properties (General)** tab, type the name (up to 32 characters) of the subclient that you want to create.
 - o Click the **Subclient Properties (Content)** tab to define the contents of the subclient.
 - o Click the **Subclient Properties (Storage Device)** tab and select a storage policy copy to associate with this subclient from the storage policy list.
 - o Click **OK** to save the subclient configuration.

MODIFYING AN AGENT, INSTANCE, OR SUBCLIENT

You can modify the configurable properties available for your agent from the agent, instance, or subclient level as per need. For subclients, typically you modify configurable properties in a user-defined subclient and not the default subclient.

It is recommended that that you do not modify the properties of an instance or subclient when a job is in progress for that specific instance or 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 be configured from the agent and instance levels.

To modify these properties:

1. From the CommCell Browser, double-click the agent, instance, or subclient you wish to modify.
Alternatively, you can also right-click the agent, instance, or subclient and then click **Properties**.
2. Navigate to the tab containing the property you wish to modify.
3. Modify the desired property and click **OK**.

OPTION	DESCRIPTION
Change the Documentum Repository User Account	<p>To perform data protection and recovery operations, the Documentum <i>iDataAgent</i> requires a user account to log on to the related Documentum docbase to access the data. The user account is the Documentum <i>Install owner</i> account that was used to install the Documentum software.</p> <p>The account must already be set up on the client.</p> <p>You can change the Documentum Repository user account details from the instance level.</p> <ol style="list-style-type: none"> 1. From the Instance Properties (General) tab, click Configure Password next to User Account. 2. In the Enter User Name and Password dialog box, type the User Name and Password. Confirm the password as well. 3. Click OK to save your settings. 4. Click Discover. The Docbase name and User Name and Password will be validated and the Storage Area and Full-text Indexes are updated. If successful, the instance is changed and the Docbase Version and Docbase Id are automatically populated. <p>Refer to User accounts and Passwords documentation for a comprehensive overview prior to using this feature.</p>
Change the Documentum Repository Instance Name	<p>You can change the instance name for a Documentum Repository at the instance level.</p> <p>To change the name of the Repository in the General tab:</p> <ol style="list-style-type: none"> 1. Enter the name (up to 32 characters) for the instance in Docbase Name. 2. Click Discover. The Docbase name and User Name and Password will be validated and the Storage Area and Full-text Indexes are updated. If successful, the instance is changed and the Docbase Version and Docbase Id are automatically populated.
Change the Documentum Repository Paths	<p>You can change the paths of the Documentum Repository at the instance level.</p> <p>To change the path names in the General tab:</p> <ol style="list-style-type: none"> 1. Type the Repository installation path in the DOCUMENTUM Path box or use Browse to locate the path. 2. Type the Repository shared path in the DOCUMENTUM_SHARED Path box or use Browse to locate the path. 3. Type Repository home path in the DM_HOME Path box or use Browse to locate the path. 4. Click Discover. The Docbase name and User Name and Password will be validated and the Storage Area and Full-text Indexes are updated. If successful, the instance is changed and the Docbase Version and Docbase Id are automatically populated.
Change the Path to DB2 Home Directory	<p>You can change the installation path of the DB2 database at the instance level.</p> <p>To change the DB2 path in the Database tab:</p> <ol style="list-style-type: none"> 1. From the Instance Properties (Database) tab, enter the path to the DB2 installation directory in the Home field. 2. Click OK to save your settings.
Change the DB2 User Account	<p>To perform data protection and recovery operations, the DB2 <i>iDataAgents</i> require a user account with one of the following privileges to access the DB2 application and database:</p> <ul style="list-style-type: none"> • Administration privileges (default) • DB2 database administration privileges • SYSADM, SYSCTRL, and SYSMANT user group privileges <p>This account must already be set up on the client. Additional accounts should be established by the DB2 database administrator. To establish additional accounts on your own, consult the appropriate DB2 application documentation.</p> <p>You can change the user account details from the instance level.</p> <ol style="list-style-type: none"> 1. From the Instance Properties (Database) tab, click Configure Password next to User Account. 2. In the Enter User Name and Password dialog box, type the User Name. 3. Click OK to save your settings. <p>Refer to User accounts and Passwords documentation for a comprehensive overview prior to using this feature.</p>
Change the Storage Area and Full-Text Index after Discovery	<p>You can change the details of the Storage Areas and Full-Text Indexes at the instance level. When Discover is selected in the General tab, the Storage Areas and Full-Text Indexes are automatically discovered and the displayed in the Storage Area and FTI tab.</p> <p>To change the Storage Area or Full-Text Index details in the Storage Area and FTI tab:</p> <ul style="list-style-type: none"> • In Storage Area Details, select a file store that was discovered and click Edit. In the Edit Operation dialog box, select the Client name from the pulldown menu.

	<ul style="list-style-type: none"> If a Full-text Index component is present, in FTI Details, select the full-text index name that was discovered and click Edit. In the Edit Operation dialog box, select the Client name from the pulldown menu and type the FTI Path. <p>Storage Areas and Full-Text Indexes can also be deleted in the Storage Area and FTI tab.</p>
Select Full-Text Index to Backup	<p>You can add, edit, or delete Full-Text Index details at the subclient level to determine which ones are backed up. Quiesce and UnQuiesce scripts can also be entered.</p> <p>To change the Full-Text Index details in the Full Text Index tab:</p> <ul style="list-style-type: none"> Select the Backup FTI checkbox to enable or disable backing up the Full-text Indexes. To add a Full-Text Index, click Add and in the Add Operation dialog box, enter the FTI Name, Client, and its FTI Path in which the full-text index resides. Click Add. To edit a Full-Text Index, select a file store and click Edit and in the Edit Operation dialog box, enter the FTI Name, Client, and its FTI Path in which the full-text index resides. To delete a Full-Text Index, select a file store and click Delete. To specify a shutdown quiesce script, type the script in Full Text Index Quiesce Script or click Browse to select it. To specify a startup quiesce script, type the script in Full Text Index UnQuiesce Script or click Browse to select it. Click OK to save the subclient configuration.
Select DB2 Database Components to Backup	<p>You can choose to backup the database in offline or online mode and specify whether to include archived logs, in the backup. Pre and Post Backup scripts can also be entered.</p> <p>To change the DB2 Database details in the Database tab at the subclient level:</p> <ul style="list-style-type: none"> Select Backup DB checkbox to enable or disable backing up the database in Online or Offline mode. To include archive logs and deleted archive logs in the backup, click Backup Archive Log and Delete Archive Log. If applicable, to specify a pre backup process, type the script in Pre Backup Process or click Browse to select it. If applicable, to specify a post backup process, type the script in Post Backup Process or click Browse to select it. Click OK to save the subclient configuration.
Select Storage Areas to Backup	<p>You can add, edit, or delete Storage Areas (file store) details at the subclient level to determine which ones are backed up. Pre and Post Backup scripts can also be entered.</p> <p>To change the Storage Areas details in the Storage Area tab:</p> <ul style="list-style-type: none"> Select the Backup SA checkbox to enable or disable backing up the Storage Areas. To add a Storage Area, click Add and in the Add Operation dialog box, enter the SA Name, Client, and its SA Path in which the Storage Area resides. Click Add. To edit a Storage Area, select a Storage Area and click Edit and in the Edit Operation dialog box, enter the SA Name, Client, and its SA Path in which the Storage Area resides. To delete a Storage Area, select a Storage Area and click Delete. To specify a pre backup process, type the script in Pre Backup Process or click Browse to select it. To specify a post backup process, type the script in Post Backup Process or click Browse to select it. Click OK to save the subclient configuration.
Create a New Index on Full Backup	<p>At the Agent level, this option determines the default behavior.</p> <p>Use this option to generate an index of the data with each full backup. This index provides quick access to the data during a browse or restore operation. If cleared, all backups will build upon the index from the previous backup.</p> <p>To create a new index on a full backup:</p> <ol style="list-style-type: none"> From the Index tab, click Create new index on full backup. Click OK to save your settings. <p>For more information about the Index and its use, see Index.</p>
Convert to Full Backup on Indexing Failures	<p>At the Agent level, this option determines the default behavior.</p> <p>Use this option to create a new index and convert the next backup to a full backup if the index is irretrievable. This index provides quick access to the data during a browse or restore operation. If cleared, the backup will fail if the index is irretrievable.</p> <p>To create a new index and convert the next backup to a full if the index is irretrievable:</p> <ol style="list-style-type: none"> From the Index tab, click Convert to full backup on indexing failures. Click OK to save your settings. <p>For more information about the Index and its use, see Index.</p>
Configure User Security	<p>You can configure user security from the agent or instance 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. Associate this object with a user group. Disassociate this object from a user group. <ol style="list-style-type: none"> From the Security tab of the associated Properties dialog box at the Agent or instance, 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.

	<p>2. Click OK to save your settings.</p>
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 Activity Control tab of the associated Properties dialog box for the instance or subclient, select or clear option(s), as desired. 2. Click OK to save your settings. <p>Refer to Activity Control documentation for a comprehensive overview prior to using this feature.</p>
View Software Version and Installed Updates	<p>The Version tab, at the Agent level displays the software version and post-release service packs and updates installed for the component.</p>
View Version of Documentum Repository	<p>Once an application has been discovered, the version of the application on which the iDataAgent was installed or upgraded on the client computer is displayed. This can be particularly useful in disaster recovery situations when you need to rebuild the environment.</p> <p>To view the version of an application,</p> <ol style="list-style-type: none"> 1. Right-click the instance and select Properties. 2. From the General tab, view the version of the application in the Version field. 3. Click OK.
Rename a Subclient	<p>You can rename an existing subclient.</p> <p>To rename a subclient,</p> <ol style="list-style-type: none"> 1. From the General tab, type the new name in the Subclient Name field. 2. Click OK to save your settings.
Change Storage Policies	<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 for command line backup and log backup from the instance level.</p> <ol style="list-style-type: none"> 1. Click Storage Device tab. 2. From the Data Storage Policy tab, select a data storage policy to associate with this subclient. 3. Click OK to save your settings.
Configure 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>To configure data transfer options,</p> <ol style="list-style-type: none"> 1. Click the Subclient Properties (Storage Device) tab, and select Data Transfer Option tab. 2. From the Data Transfer Option tab, do the following: <ul style="list-style-type: none"> ○ Choose the appropriate compression option for this subclient. ○ Select Throttle Network Bandwidth to set the network bandwidth settings. 3. Click OK
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> <p>From the Data Storage Policy tab, click Data Paths to view the data paths used by the subclient to access the storage media for data protection operations.</p>
Set the Number of Streams	<p>You can set the number of streams for database, log, or Storage Area backup at the subclient level.</p> <ol style="list-style-type: none"> 1. From the Storage Device tab: <ul style="list-style-type: none"> ○ Select the Streams tab. ○ Select the number of data streams for the Database Streams field. ○ Select the number of log streams for the Log Streams field. ○ Select the number of Storage Areas streams by selecting a Storage Area and clicking Edit. In the Edit Operation dialog box, use the spin boxes to set the number of streams. ○ Select Data Threshold and enter a number to specify whether to run the backup job immediately when the total number of reserved streams is greater than or equal to the value in this field. 2. Click OK to save your settings. <p>Refer to Streams documentation for a comprehensive overview prior to using this feature.</p>
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. You can also change the user account for running Pre/Post processes for data protection operations of agents residing on Windows platforms.</p> <p>Consider the following before modifying the Pre/Post processes for the subclient.</p> <ul style="list-style-type: none"> • We recommend not configuring a pre/post process for a subclient that is currently running a data protection or archive operation.

	<ul style="list-style-type: none"> • Verify that there are no pre/post processes already assigned for the subclient. • Pre-process commands for the iDataAgents will be executed only when the necessary resources (e.g., media, library, drive, etc.) are available. <ol style="list-style-type: none"> 1. Click the Subclient Properties (Pre/Post Process) tab. 2. Click one of the following phases and type the full path of the process that you want executed during that phase. Alternatively, click Browse to locate the process (applicable only for paths that do not contain any spaces). <ul style="list-style-type: none"> ○ PreBackup ○ PostBackup Click OK. 3. If you want to run a post backup process for all attempts to run that job phase, then select the Run Post Backup Process for all attempts checkbox. 4. For subclients on Windows platforms, if Run As displays Not Selected, or if you want to change the account that has permission to run these commands, click Change. <ol style="list-style-type: none"> a. In the User Account dialog box, select Use Local System Account, or select Impersonate User and enter a user name and password. Click OK. b. If you selected Local System Account, click OK to the message advising you that commands using this account have rights to access all data on the client computer. <p>Refer to Pre/Post Processes documentation for a comprehensive overview prior to using this feature.</p>
Enable/Disable Data Encryption	<p>You can enable data encryption from the instance and subclient level. When accessing the Encryption tab from the instance level, you are selecting either Network and/or Media-side encryption for third-party Command Line operations. When accessing this same tab from the subclient level, the settings selected apply only to the selected subclient content for operations run from the CommCell Console.</p> <p>Consider the following before enabling encryption:</p> <ul style="list-style-type: none"> • Encryption must be enabled at the client level prior to configuring any subclients residing on that client. • If you are attempting to configure for third-party Command Line operations, enable data encryption from the instance level. <p>Encryption can be set from the Encryption tab of the associated Properties dialog box at the instance and subclient level.</p> <p>Refer to the Data Encryption documentation for a comprehensive overview prior to using this feature</p>

DELETING AN INSTANCE OR SUBCLIENT

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

When you delete an instance or subclient, 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.

DELETE 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.

DELETE 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, right-click the user-defined subclient that you want to delete, and then click **Delete** from the shortcut menu.
 2. A confirmation message is displayed, asking if you want to delete the subclient.

Click **No** to cancel the deletion and retain the subclient, or click **Yes** to continue the deletion.

[Back to Top](#)

Backup - Documentum iDataAgent (DB2)

Basic | **Advanced**

TABLE OF CONTENTS

Overview

What Gets Backed Up

Your first Full Backup

- Running a Full Backup Immediately
- Scheduling a Full Backup

Viewing Backup Job History

OVERVIEW

Once you have created an instance, you can run your first backup operation on the default subclient. The first backup on the default subclient will be a full backup and will include the entire Documentum Repository including the DB2 database (data and logs).

Backups can be run on demand or can be scheduled to be run at a later point of time. The following section explains in detail the procedure to run your first backup operation.

WHAT GETS BACKED UP

The Documentum iDataAgent backs up the following types of data:

- DB2 database files, log files, and History files.
- Storage Areas
- Full-text Indexes (optional)

However, to backup DB2 application files associated with the DB2 installation, you will need to use the appropriate File System iDataAgent.

YOUR FIRST FULL BACKUP

Backups for any client start with a full backup. The full backup becomes a baseline to which subsequent backup types are applied.

A full backup contains all the data that comprises a subclient's contents. If a client computer has multiple agents installed, then the subclients of each agent require a full backup in order to secure all of the data on that client.

- Before running the first backup, ensure to configure the Storage Policy and the streams that are required for backing up the Storage Areas and Database.
- Prior to running your first backup, in a Single-Server Environment where the Documentum iDataAgent backs up all the components (database, storage areas, full-text indexes), you must run `Db2_install.sh` script in the `<installation_path>/iDataAgent` folder to ensure proper library linking.
- DB2 database configuration parameters need to be enabled before you backup DB2 database files. See [Enable Database Configuration Parameters](#) for more information.


The following sections provide step-by-step instructions for running your first backup, either immediately or at a scheduled time.

RUNNING A FULL BACKUP IMMEDIATELY:

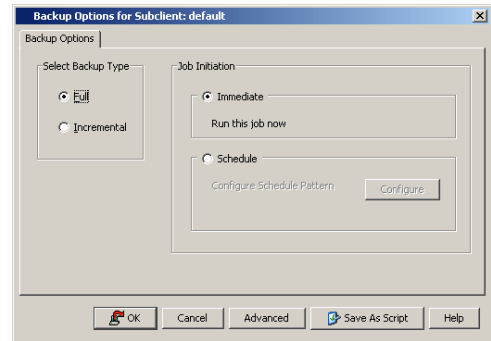
1. From the CommCell Browser, right-click the default subclient, and click **Backup**.

If a backup is run with the default subclient, all data types (database, storage areas, full-text indexes) are automatically selected and will be backed up.

2. From the **Backup Options** dialog box, select **Full** backup type.




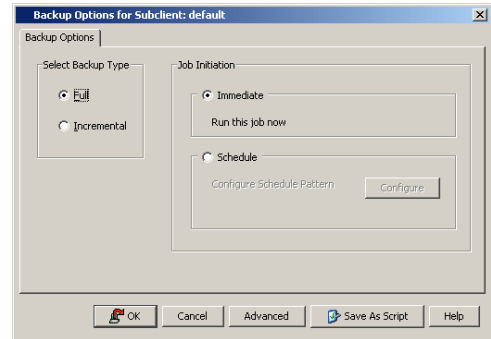
3. Select **Immediate** to run the backup operation immediately.



4. Click **OK**.

You can track the progress of the backup job from the **Job Controller** window. If you are using a stand-alone drive, you are prompted to load a specific cartridge into the drive. If you are using a library, you will not receive this prompt. The system loads the tapes automatically. Your cartridges should be appropriately labeled. This will enable you to locate the correct cartridge for a restore job, if necessary. When the backup has completed, the Job Controller displays **Completed**.

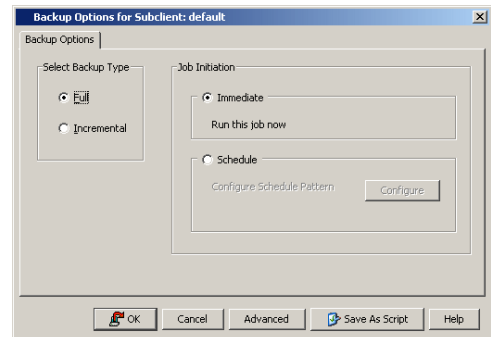

Once completed, the details of the job can be viewed in the **Job History** window.



SCHEDULING A FULL BACKUP:

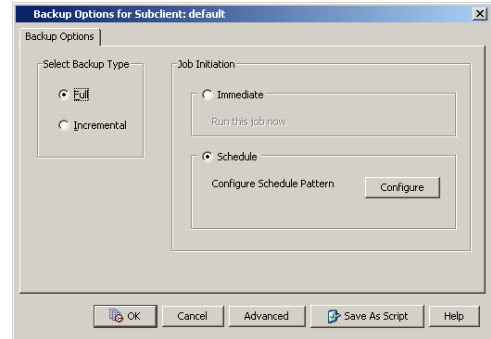
If the database and storage areas are backed up with separate subclients, ensure to schedule the backups so that the database backup is performed first followed by the storage area and full-text index backups.

1. From the CommCell Browser, right-click the subclient, and click **Backup**.
2. From the **Backup Options** dialog box, select **Full** backup type



3. Select **Schedule** to schedule the backup for a specific time.

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




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

Specify the following options:

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

You can also confirm and/or edit (where applicable) your choices from the **Job Summary** tab. For a monthly or yearly schedule, you can select either a standard calendar or a custom calendar.


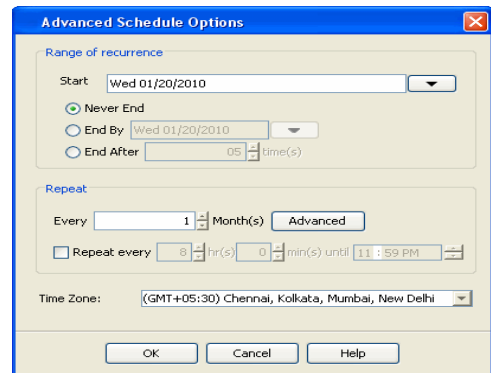
Click **OK** once all the desired options have been configured.



6. Click **OK**.

Your backup operation will execute according to the specified schedule. Once initiated, you can track the progress of the backup job from the **Job Controller** window.

Once completed, the details of the job can be viewed in the **Job History** window.



VIEWING BACKUP JOB HISTORY

Once a backup job has completed, it may be useful to view specific historical information about the job, such as the following:

- Items that failed during the job
- Items that succeeded during the job
- Details of the job
- Media associated with the job
- Events of the job
- Log files of the job.

This information can be viewed in the Job History window, which is accessible at the client computer, iDataAgent, instance, or subclient level in the CommCell Console.

The history provided depends on the entity from which the job history is viewed. For example, viewing job history from the subclient level will yield information for the jobs performed for that subclient. Conversely, viewing job history from the instance level will yield information for jobs run from all subclients within that instance.

To view the backup job history:

1. From the CommCell Browser, right-click the entity (client computer, iDataAgent, instance, or subclient) whose backup history you want to view, click **View**, and then click **View Backup History**.

2. From the Backup History filter window select the filter options, if any, that you want to apply, and then click OK. The system displays the Backup Job History window.
3. Once you have chosen your filter options, they are displayed in the **Backup Job History** window.
To view the additional options discussed above, right-click the desired job choose the appropriate option.
4. Click **OK**.

[Back to Top](#)

Backup - Documentum iDataAgent (DB2)

Basic | Advanced

TABLE OF CONTENTS

Overview

What Happens During Documentum Backup

Incremental backups

Starting an Incremental Backup

Database Backups

Offline Backups

Online Backups

Archive Log Backups

Selective Copy - Full Backups

Enable Database Configuration Parameters

Storage Area Backups

Full-text Index Backups

Using Documentum X-plore as Indexing Software

Managing Jobs

Restarting Jobs

Controlling Jobs

Additional Options

OVERVIEW

Once you have performed a full backup of the Documentum Repository initially, the subsequent backups can be either Full backups or Incremental backups.

A full backup contains all the data that comprises a subclient. Though, full backups provide the most comprehensive protection of data, it consumes the most amount of time and resources. A full backup is supported for backing up database files, Storage Areas, and Full-text Indexes. Optionally, log files can be backed up as well.

To streamline the backup process and to optimize the time and resource, you can run incremental backups. You can perform these backups from the CommCell Console. Incremental backups are supported for backing up database files and Storage Areas.

Backups can be run on demand or can be scheduled to be run at a later point of time. However, it is always recommended that you schedule your backups at frequent intervals. For example, full backups can be scheduled on a weekly basis and incremental backups can be scheduled on a daily basis.

The following sections describes the various types of backups supported by this agent.

WHAT HAPPENS DURING DOCUMENTUM BACKUPS

In order to achieve a consistent backup of Documentum components, Bull performs the following during a repository backup:

1. Pre-backup phase includes:
 - a. Quiesces the Full-Text index (indexer process).
 - b. Inactivate `dm_DMclean` and `dm_DMfilesan` jobs.
 - c. Create Bull Index.
2. Backs up the Full-Text Index (FTI).
3. Backs up the database.
4. Backs up the log files.
5. Backs up the storage area.
6. Post-backup phase includes:
 - a. Unquiesces the Full-Text Index (indexer process).
 - b. Activates `dm_DMclean` and `dm_DMfilesan` jobs.
7. Archives Bull Index and copies it to media.

INCREMENTAL BACKUPS

An incremental backup contains only data that is new or has changed since the last backup, regardless of the type. On average, incremental backups consume

RELATED TOPICS

Job Management


Provides comprehensive information on managing jobs.

far less media and place less of a burden on resources than full backups.

Always ensure that your first backup is a full backup that will contain all the data that comprises a subclient's contents.

The illustration on the right clarifies the nature of full and incremental backups. For simplicity, assume there is a file system that contains six files as represented in the figure.

Backup #1 is a full backup and therefore writes all the data, changed and unchanged, to the backup media. Backups #2 through #n-1 are incrementals and only back up those files that have changed since the time of the last backup, regardless of the type. For example, files A, B, and E changed after the full backup and were therefore backed up in Backup #2. Backup #4 backed up files A and D because both files were modified sometime after Backup #3 occurred. File F did not change; hence, it was not backed up in any of the incremental backups, but it was included in both full backups, which, by definition, back up everything.



By default, during the database backup phase of an incremental backup, only logs are backed up. This can be overridden by setting the `SINCREMENTALDCTMDBBACKUP` registry key to `Y`.

When a Storage Area is backed up without any new documents, the backup job completes but nothing is backed up.

STARTING AN INCREMENTAL BACKUP

- From the CommCell Browser, expand **Client Computers** by double-clicking **Client Computers | Documentum | <instance name>**. The default and other subclients (if available) are displayed on the right-hand windowpane.
- To back up the default subclient, right-click the subclient, and click **Backup**.
- From the **Backup Options** dialog box, select from the following options as desired:

Backup Type:

 - Select **Incremental**.

Backup Schedule:

 - Select **Immediate** if you want to run the backup now.
 - Select **Schedule** if you want to schedule the backup for a specific time.
- Click **OK**. You can track the progress of the backup job from the **Job Controller** window. If you are using a stand-alone drive, you are prompted to load a specific cartridge into the drive. If you are using a library, you will not receive this prompt. The system loads the tapes automatically. Your cartridges should be appropriately labeled. This will enable you to locate the correct cartridge for a restore job, if necessary.
- When the backup has completed, the Job Controller displays **Completed**.

DATABASE BACKUPS

The Documentum iDataAgent enables you to backup the database in different environments based on need. Database files can be backed up in full or incremental backups. These are described in the following sections.

If the database and storage areas are backed up with separate subclients, ensure to schedule the backups so that the database backup is performed first followed by the storage area and full-text index backups.

OFFLINE BACKUPS

Offline backups are performed when the database is shut down and unavailable for use. Offline backups are always full backups and do not include the archived log files.

In order to backup the database when it is offline, you can create a separate user-defined subclient for offline backup. See [Creating Subclient to Backup Offline Databases](#) for step-by-step instructions.

When performing Documentum offline backups from the CommCell Console, shutdown the Documentum Repository services before performing the backup. This can also be run as a Pre Process during backing up to shut it down gracefully. A Post Post backup script can then be run to start it up again after running the backup.

ONLINE BACKUPS

Online backups are performed when the database is online. Data files are included when you back up the database in online mode. You can perform a full or incremental backup when the database is online. These backups do not require database downtime and are extremely useful when you need to perform a point-in-time restore of the database.

In order to backup the database when it is online, you can create a separate user-defined subclient for online backup. See [Creating Subclient to Backup Online Databases](#) for step-by-step instructions.

ARCHIVE LOG BACKUPS

Archive log backups are useful when you want to recover database transactions that have been lost due to an operating system or disk failure. You can apply these archive logs to an online backup in order to recover a database.

Full backups generally includes both the data and logs. However, you can also take backups of the archive logs only. In order to backup only the archive logs, you need to create a separate subclient for archive logs. See [Creating Subclient to Backup Archive Log Files](#) for step-by-step instructions.

By default, the first backup operation on the default subclient will include the entire database (data and archive logs) when the database is online.

Frequent backups of database log files should be performed. Without them, a database can be recovered only to the time of its last full backup.

SELECTIVE COPY - FULL BACKUPS

A selective copy allows you to copy backup data selectively from a source copy to this copy, providing for better tape rotation. Since only selective backups can be copied to selective copies, the selective copies cannot be promoted to the primary copy, only synchronous copies can be promoted. Note that the data selection process does not have to be the same for all auxiliary copies.

During an auxiliary copy operation or a data protection operation that creates an inline copy, only those backups from the primary copy that meet certain criteria will be copied to a selective copy. You can define a selective copy to be time-based, automatically selected, or automatically not selected on the primary copy.

If the copy is defined as All Fulls, all full backups on the primary copy will be copied during an auxiliary copy operation or a data protection operation that creates an inline copy. If the copy is defined as time-based, only the first or last full backup that occurs within each selected weekly, monthly, quarterly, half-yearly, or yearly interval will be copied. You have the option of creating and associating a custom calendar to the copy, so that the intervals can be further customized.

Selective copy will not only copy Selective Online Full and Selective Offline Full jobs, it will also copy regular Full jobs with the linked Log backup jobs.

You need to create a separate subclient to perform a selective copy full backup of the database. See [Configuring Subclient for Selective Copy with Full Backups](#) for step-by-step instructions.

Selective Copy is applicable only for jobs that have database backed up as part of Full jobs.

ENABLE DATABASE CONFIGURATION PARAMETERS

Prior to your first backup, you need to enable DB2 database configuration parameters as follows.

BACKUPS IN GENERAL

To back up any type of DB2 data online or offline, update the database configuration parameters when DB2 agent is installed on a cluster or multiple instances of Calypso are installed.

- From the DB2 console, type the following command to set the **LOGARCHOPT1** parameter:

```
db2 update db cfg for <database name> using LOGARCHOPT1''CvClientName=<CvClientName>,CvInstanceName=<CvInstanceName>''
```

Example:

```
db2 update db cfg for test_db using LOGARCHOPT1''CvClientName=testhost,CvInstanceName=Instance001''
```

- Type the following command to set the **VENDOROPT** parameter:

```
db2 update db cfg for <database name> using VENDOROPT''CvClientName=<CvClientName>,CvInstanceName=<CvInstanceName>''
```

Example:

```
db2 update db cfg for test_db using VENDOROPT''CvClientName=testhost,CvInstanceName=Instance001''
```

- Type the command to set the **LOGARCHMETH2** to **OFF**

```
db2 update db cfg for <database name> using LOGARCHMETH2 OFF
```

where:

<CvClientName> identifies the name of the computer where the DB2 iDataAgent is installed

<CvInstanceName> identifies the name of the Calypso instance (e.g., **Instance001**). This should be the name of the DB2 iDataAgent instance associated with the appropriate CommServe.

LOG BACKUPS

To back up DB2 log files for DB2 version 8.2 or later using the SBT interface, enter the following command to set the required parameters:

```
db2 update database configuration for <database name> using LOGARCHMETH1 VENDOR:<install path><shared library file>
```

where:

`<install path>` identifies the base install directory

`<shared library file>` identifies the name of the shared library file `libDB2Sbt.*` (where `*` identifies the file extension, which will vary per the operating system)

With the SBT interface, log files are staged to an archive path, and an On Demand backup of the log files is performed when a certain threshold is reached. Alternatively, the user can submit a backup request for any number of log files from the CommCell Console.

The SBT interface program takes the logs from the DB2 directory and transfers them first to the threshold directory that was defined during the DB2 iDataAgent install and then to tape. If you subsequently use the Delete option, the logs are deleted from the threshold directory; otherwise, the logs are kept in this directory. By default, the threshold directory can hold 1000 logs. If the threshold directory attempts to hold more than 1000 logs, a command line backup for the logs starts automatically, and the logs are transferred to tape.

You can control the threshold value by adding the `sDb2ThresholdALFN <number>` registry key to the Db2Agent registry, where `<number>` indicates the number of logs you want to hold in the t

ONLINE BACKUPS

To back up any type of DB2 data online, enable the `LOGRETAIN` database configuration parameter by entering the following command:

```
db2 update database configuration for <database name> using LOGRETAIN on
```

INCREMENTAL BACKUPS

To run delta or incremental DB2 backups, enable the `TRACKMOD` database configuration parameter by entering the following command:

```
db2 update database configuration for <database name> using TRACKMOD on
```

Whenever you change the `TRACKMOD` database configuration parameter, be sure to run a full offline database backup. Otherwise, the `TRACKMOD` parameter will not take effect.

STORAGE AREA BACKUPS

The Documentum iDataAgent enables you to back up both local and distributed Storage Areas that are distributed across different computers. With distributed storage areas, the File System iDataAgent should be installed on the client computers that are hosting the Storage Areas.

Both full and incremental backup types are supported for Storage Areas. The first backup of a Storage Area is performed as a Full backup. Storage Areas can be backed up in a full Repository backup or as a separate backup. For the latter case, see Create Subclient to Only Contain Storage Areas for step-by-step instructions.

A pre or post backup process can be run when backing up Storage Areas.

In order to discover far stores and their associated clients in a distributed environment, ensure that they are configured correctly.

FULL-TEXT INDEX BACKUPS

The Documentum iDataAgent enables you to back up Full-text Indexes. The File System iDataAgent should be installed on the client computers that are hosting the Full-text Indexes.

Full-text Indexes can be backed up only in full backups. Even if an incremental backup job is run, Full-text Indexes will always be backed up as a full backup. Full-text Indexes can be backed up in a full Repository backup or as a separate backup. For the latter case, see Create Subclient to Only Contain Full-text Indexes for step-by-step instructions.

When backing up a Full-Text Index, it necessary to quiesce the Index Server to ensure consistent backups. For this purpose, quiesce and unquiesce scripts are required. These are used to quiesce and unquiesce the Indexer process in the Index Server during the backup operation respectively.

For Unix based Documentum iDataAgents, these two sample scripts are provided in the `<installation_folder>/Base` folder on the computer where the Documentum iDataAgent has been installed.

For windows based Documentum iDataAgents, the FTI quiesce and unquiesce scripts are similar to the scripts in the following example:

```
cvftishutdown.bat
-----
E:\Documentum\fulltext\IndexServer\bin\nctrl.exe stop

cvftistartup.bat
-----
C:\WINDOWS\system32\net.exe start FASTDSService
E:\Documentum\fulltext\IndexServer\bin\nctrl.exe start
```

USING DOCUMENTUM X-PLORE AS INDEXING SOFTWARE

For Windows based Documentum /DataAgents, the Documentum X-plore indexing server quiesce and unquiesce scripts are similar to the scripts in the following example.

These scripts have to be modified prior to backup to suit your environment:

CVXPLORESHUTDOWN.BAT

```
@REM set GALAXY_HOME=C:\TEMP

@REM set GALAXY_HOME=C:\TEMP

@REM set GALAXY_LOG=%GALAXY_HOME%

set CV_XPLORE_LOG=C:\TEMP\CV_xPlore.log

set errorlevel=0

@REM CALL Index Agent Unquiesce

echo "Index Agent quiesce begin at %date% %time%..." >> %CV_XPLORE_LOG%

@@ "C:\xPlore\dsearch\admin\xplore.bat suspendDiskWrites()"

set PH_1=%errorlevel%

IF NOT %PH_1% EQU 0 (

echo Index Agent quiesce Failed reason >> %CV_XPLORE_LOG%

GOTO FAIL

) ELSE (

echo "Index Agent quiesce is done at %date% %time%..." >> %CV_XPLORE_LOG%

)

:PASS

echo Index Agent quiesce or DB suspend Passed.

exit /b %errorlevel%

:FAIL

echo Echo Index Agent quiesce or DB suspend Failed.

exit /b %errorlevel%
```

CVXPLORESTARTUP.BAT

```
@REM set GALAXY_HOME=C:\TEMP

@REM set GALAXY_HOME=C:\TEMP

@REM set GALAXY_LOG=%GALAXY_HOME%

set CV_XPLORE_LOG=C:\TEMP\CV_xPlore.log

set errorlevel=0

@REM CALL Index Agent Unquiesce

echo "Index Agent Unquiesce begin at %date% %time%..." >> %CV_XPLORE_LOG%
```

```

@@ "C:\xPlore\dsearch\admin\xplore.bat resumeDiskWrites()"

set PH_1=%errorlevel%

IF NOT %PH_1% EQU 0 (

echo Index Agent Unquiesce Failed reason >> %CV_XPLORE_LOG%

GOTO FAIL

) ELSE (

echo "Index Agent Unquiesce is done at %date% %time%..." >> %CV_XPLORE_LOG%

)

:PASS

echo Index Agent Unquiesce or DB Resume Passed.

exit /b %errorlevel%

:FAIL

echo Either Index Agent Unquiesce or DB Resume Failed.

exit /b %errorlevel%

```

MANAGING JOBS

Once you initiate the backup operation, a backup job is generated in the Job Controller. 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. This configuration can be changed at any time; however, changes to this configuration will affect all jobs run in the entire CommCell.

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. In the CommCell Console, click the **Control Panel** icon, then double-click **Job Management**.
2. In the **Job Restarts** tab of the Job Management dialog box, select a job type from the **Job Type** list, and then select the **Restartable** check box.
3. To change the maximum number of times the Job Manager will try to restart a job, select a number in the **Max Restarts** box.
4. To change the time interval between attempts by the Job Manager to restart the job, select a number of minutes in the **Restart Interval (Mins)** box.
5. Click **OK** to save your changes.

CONFIGURE JOB RESTARTABILITY FOR AN INDIVIDUAL JOB

1. From the **Backup Options** dialog box, click **Advanced**, then select the **Job Retry** tab and specify the following as desired:
 - o **Total Running Time** - The maximum elapsed time, in hours and minutes, from the time that the job is created.
 - o **Number of Retries** - The number of times that Job Manager will attempt to restart the job.
 - o **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".
2. 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. Only preemptible jobs can be suspended.
RESUME	Resumes a job and returns the status to Waiting, Pending, Queued, or Running depending on the availability of resources or the state of the operation windows and activity control settings. Backup jobs for this Agent are resumed from the point-of-failure.
KILL	Terminates a job.

These controls can be applied to:

- All jobs in the Job Controller.
- All selected jobs in the Job Controller providing you have the correct security associations at the proper level for each job selected.
- All data protection operations running for a particular client or client/agent.
- All data protection operations running for a particular MediaAgent.

1. From the Job Controller of the CommCell Console, right-click the job and select **Kill**, **Suspend**, or **Resume** as desired.

o When killing a job:

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.

o When suspending a job:

The job status may change to **Suspend Pending** for a few moments while the operation completes. The job status then changes to **Suspended**.

o When resuming a job:

As the Job Manager attempts to restart the job, the job status changes to **Waiting**, **Pending**, or **Running**.

- When backup jobs are suspended before the chunk size is committed, the job will restart the backup from the uncommitted chunk. Therefore, the backup size of a suspended job may be larger than the size of a normal backup job.
- When a backup job is killed, you need to restart these Documentum jobs manually using Documentum software:
 - o dm_DMCClean
 - o dm_DMFileScan

ADDITIONAL OPTIONS

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

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

OPTION	DESCRIPTION
Startup Options	<p>The start options are accessible from the Startup tab in the Advanced Backup Options window.</p> <p>These options facilitate the startup behavior of the backup job.</p> <p>The Default Priority option dictates the type of job that will be used in determining how the Job Manager will allocate resources for this job.</p> <p>The Change Priority option specifies 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.</p> <p>The 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.</p> <p>Refer to the Job Priority and Priority Precedence documentation for a comprehensive overview prior to using these options.</p>
Job Retry Options	<p>The job retry options are accessible from the Job Retry tab in the Advanced Backup Options window.</p> <p>These options enable you to configure the retry behavior of the backup jobs.</p> <ul style="list-style-type: none"> • 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. • 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

	<p>Time has elapsed, even if its state is "Running". This option is available only if you have specified a Total Running Time.</p> <p>Refer to the Job Management documentation for a comprehensive overview prior to using these features.</p>
Start New Media	<p>This option is accessible from the Media tab in the Advanced Backup Options window.</p> <p>This option starts the backup/archive operation on a new media, which causes the following to occur:</p> <ul style="list-style-type: none"> • If removable media is used, the current active media is marked as Appendable and a new media is used for the backup/archive. • If disk media is used, a new volume folder is created for the backup/archive. <p>If cleared, the operation automatically uses the current active media or volume.</p> <p>This media management feature provides a degree of control over where the data physically resides, for example helping you to meet security or performance goals. This feature is useful in situations where you would like the data to reside on a new media, not appended to a media already containing backup/archive data.</p> <p>Refer to the Creating an Exportable Media Set section of the Export Media documentation for a comprehensive overview prior to using this feature.</p> <p>Another form of the Start New Media option is available from the library properties. Refer to the Start New Media section of the Library Properties documentation for a comprehensive overview.</p>
Mark Media Full	<p>This option is accessible from the Media tab in the Advanced Backup Options window.</p> <p>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.)</p> <p>This media management feature provides a degree of control over where the data physically resides, for example helping you to meet security or performance goals. This feature prevents any other data from being written 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.)</p> <p>Refer to the Create an Exportable Media Set section of the Export Media documentation for a comprehensive overview prior to using this feature.</p> <p>Refer to the Export Media documentation for a comprehensive overview prior to using this feature.</p>
Allow other Schedules to use Media Set	<p>This option is accessible from the Media tab in the Advanced Backup Options window.</p> <p>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:</p> <ul style="list-style-type: none"> • 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. <p>Refer to the Creating an Exportable Media Set section of the Export Media documentation for a comprehensive overview prior to using this feature.</p>
Data Path Options	<p>The data path options can be accessed from the Data Path tab in the Advanced Backup Options window.</p> <p>Data can be backed up to any compatible library and drive type in the CommCell. By default the system automatically identifies the data path for backup operations as configured in the CommCell. You can change this data path at the backup level by selecting a specific MediaAgent, library or drive pool in the Advanced Backup Options dialog. The data path change can be specified at the backup, schedule and the schedule policy levels. Note that the data path changes specified are only applicable only for the specific backup job/backup schedule.</p> <p>Ensure that the specified MediaAgent and the library or drive pool are available and accessible, failing which the job will go into Pending state; failover protection does not apply here. Also, note that this setting overrides all other settings, including the data path settings at the Storage Policy, Subclient and the alternate data path (Gridstor).</p> <p>The change data path options are useful in the following situations:</p> <ul style="list-style-type: none"> • The preferred MediaAgent, library or drive pool is not available during the backup. • You wish to use a different MediaAgent, library or drive pool for load balancing purposes. • You wish to direct that backup to a particular location for compliance purposes. <p>Refer to Change Data Path documentation for a comprehensive overview prior to using this feature.</p>
Vault Tracker	<p>You can access the vault tracking options from the Vault Tracker tab in the Advanced Backup Options window.</p> <p>Several Vault Tracking options are available to export and track media. These options are available only when a Vault Tracker license is available in the CommServe, and 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.)</p> <p>Refer to the following documentation for a comprehensive overview prior to using this feature:</p> <ul style="list-style-type: none"> • VaultTracker if a standard CommServe license is available. • VaultTracker Enterprise if the VaultTracker Enterprise license is available.
Alerts	<p>You can configure alerts for the specific backup job from the Alert tab in the Advanced Backup Options window.</p> <p>Refer to Alerts documentation for comprehensive overview prior to using this feature.</p>

Back to Top

Restore - Documentum iDataAgent (DB2)

Basic Advanced

TABLE OF CONTENTS

Overview

Your First Restore

Documentum File Browse and Restore
Component-Level Browse and Restore

Viewing Restore Job History

Viewing Restored Items Associated With a Restore Job

OVERVIEW

As restoring your backup data is very crucial, it is recommended that you perform a simple restore operation immediately after your first full backup to understand the process. The following sections explain the steps for performing simple browse and restore operations.

Prior to running your first restore operation, ensure the following:

- Verify that the CommServe, MediaAgent and media library are powered on.
- You have successfully done a backup of the data to be restored.
- Verify that the media from which you wish to perform the restore is available in the CommCell. Use the List Media feature to identify the media required by the restore operation.
- Always ensure that sufficient space has been allocated for the restore.

Restores can be run on-demand or can be scheduled to be run at a later point of time. The following sections explain the procedure to perform your first restore operation.

YOUR FIRST RESTORE

The restore operation can be done using one of the following methods:

- File-level Browse and Restore, which is useful when you want to restore documents.
- Component-level Browse and Restore, when you want to restore components of a Documentum Repository (database, storage area, or full-text indexes).

For this first simple restore, let us perform an in-place restore of documents and components of a Documentum Repository.



DOCUMENTUM FILE BROWSE AND RESTORE

This type of Browse and Restore is geared towards end-users and illustrates performing a basic browse and restore of documents:

1. From the CommCell Browser, to browse and restore documents, select one of the following methods:

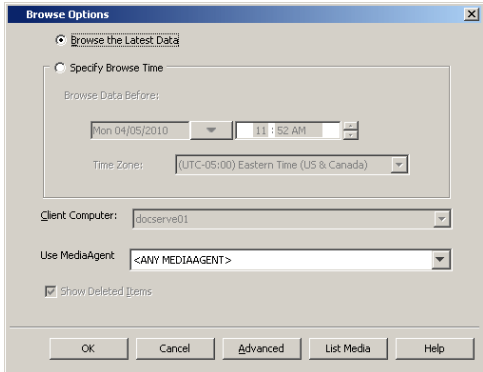

Instance Level:

- Right-click a Documentum instance, select **All Tasks** then click **Browse Backup Data**.
- From the **Browse Options** dialog box, accept the default **Browse the Latest Data**.
- Then click **Documentum View** option and specify a subclient (e.g., default) that contains documents.
- Click **OK**.



Subclient Level:


- Right-click a Documentum subclient that contains documents and select **Documentum File Browse**.
- From the **Browse Options** dialog box, accept the default **Browse the Latest Data**.
- Click **OK**.



2. The **Browse** window is displayed, containing data objects such as files, folders, etc. 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 tree.

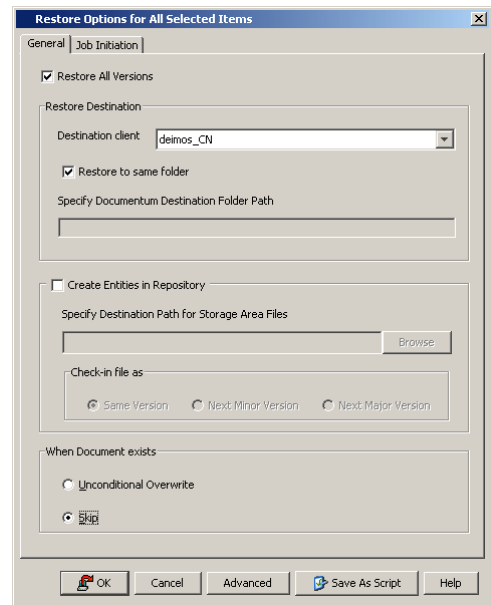
Note that the window displays only the data that was obtained by a data protection operation. Data that is excluded by a filter, or data which did not exist as of the specified browse time does not appear in the window.



3. Select the instance node in the left pane. The files will be selected in the right pane. Select one or more files and click the **Recover All Selected** button at the bottom of the window.

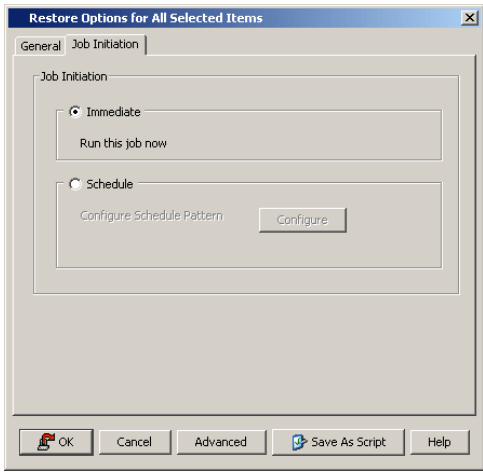
4. From the **Restore Options (General)** tab:

- Select the name of the source client computer in the **Destination Client** from the list.
- If performing an in-place restore, click **Restore to same folder**. Otherwise, specify a path in **Specify Documentum Destination Folder Path**.
- Click **Create Entities in Repository** to specify a destination storage area and check-in either a current or next major/minor versions of documentum files. The Documentum files from the MediaAgent will be dumped on to a local staging path and then restored on the destination computer.
 - In **Specify Destination Path for Storage Area Files**, enter the file system path on the destination computer to which the files will be restored. Typically, this would be a system path on the Primary Content Server. Click one of the following:
 - **Current Active Version** to check the file in with the same version number that is currently active in Documentum.
 - **Next Minor Version** to check the restored file to the next Minor Version. For example, if the restored document was a Minor version (e.g., 1.2), it will be check it into the next Minor version (e.g., 1.3). If the restored document is a Major version (e.g., 2.0), it will be checked into the next Minor version (e.g., 2.1).
 - **Next Major Version** to check-in the file to the next major version. For example, if the restored document was a Major version (e.g., 2.0), it will check it into the next Major version (e.g., 3.0). If the restored document was a Minor version (e.g., 1.2), it will check it into the next Major version (e.g., 2.0).
- Select **Skip** to restore and recover the documents.




5. If you need to run the backup operation immediately, select **Immediate** from the **Job Initiation** tab.


Go to step 11.



- 6. If you need to schedule the restore operation, select **Schedule** from the Job Initiation tab and click **Configure**.
- 7. From the **Schedule Details** dialog box that appears, select the appropriate scheduling options.
Click **Options** to view the **Advanced Schedule Options** dialog box.



- 8. From the **Advanced Schedule Options** dialog box, select a specific time zone from which the job schedule time will be based.
Click **OK** to close the **Advanced Schedule Options** dialog box.



- 9. Click **OK** to close the **Schedule Details** window.
- 10. Click **OK** to close the **Restore Options** window.
You can track the progress of the restore job from the **Job Controller** window.
Once completed, the details of the job can be viewed in the **Job History** window.

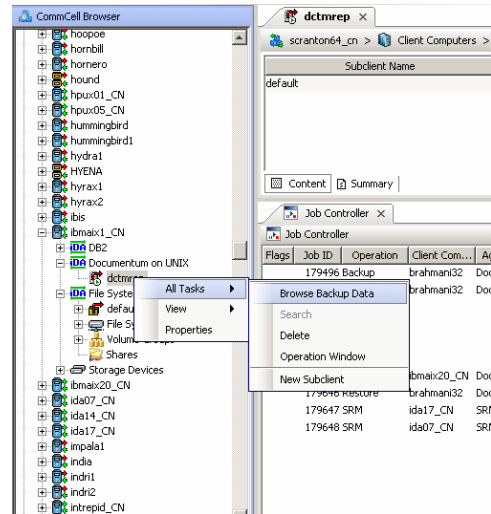
COMPONENT-LEVEL BROWSE AND RESTORE

This type of Browse and Restore is geared towards administrators and illustrates performing a basic browse and restore of all components of a Documentum Repository (database, storage area, or full-text indexes):

To ensure consistent restores, it is recommended to shutdown the Documentum Repository and the database in an appropriate mode before restoring.

- 1. From the CommCell Browser, to browse and restore components of a Documentum Repository, select one of the following methods
 - Right-click the Documentum agent, select **All Tasks** then click **Browse Documentum**.
 - Right-click a Documentum instance, select **All Tasks** then click **Browse Backup Data**.

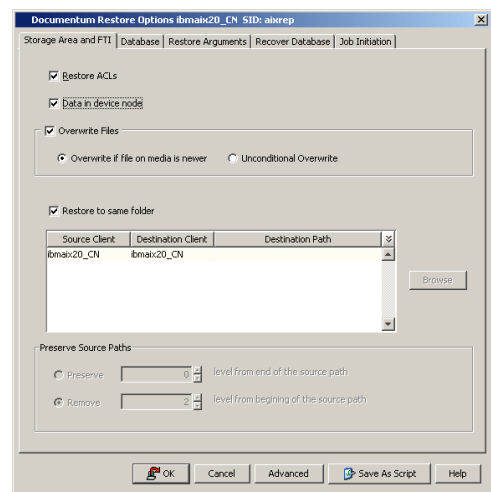


- From the **Browse Options** dialog box, click **OK** to execute the browse using the **Browse the Latest Data** option.



- The **Browse** window is displayed, containing data objects such as database, storage area, full-text index 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 tree.

Note that the window displays only the data that was obtained by a data protection operation. 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.

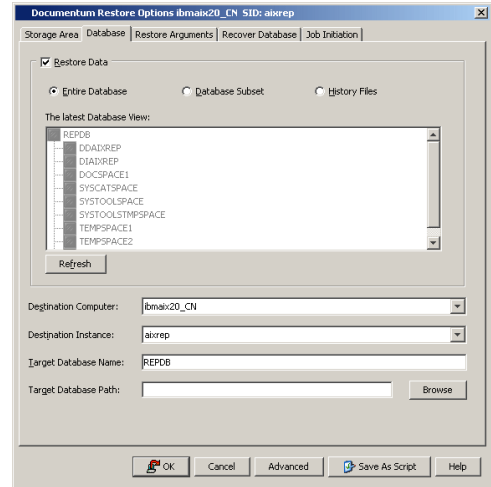
- Select the components in the right pane that are required to restore. Click the **Recover All Selected** button at the bottom of the window.
- To restore the storage areas and full-text indexes to the same folder, from the **Documentum Restore Options (Storage Area and FTI)** tab:
 - Select **Overwrite Files** to overwrite all backed up storage areas.
 - Select **Restore to same folder** to restore the storage area to the same folder.



- To restore the database, from the **Documentum Restore Options (Database)** tab:
 - Select **Restore Data** and select **Entire Database**.

- Select the name of the client computer in the **Destination Client** from the pulldown menu.
- Select the name of the DB2 instance in the **Destination Instance** from the pulldown menu.
- Type the database name of the **Target Database Name** box.
- Type the database path of the **Target Database Path** box.

Ensure to include in this field just the letter of the appropriate drive followed by a colon (e.g., I:). Otherwise, the target database will not be accessed.



7. To restore the backup images available for restore, from the **Documentum Restore Options (Restore Arguments)** tab:

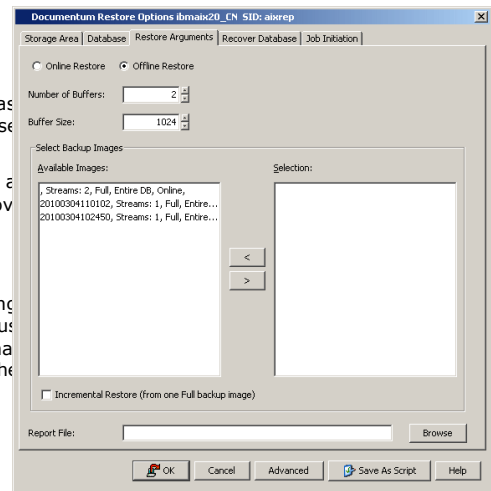
- Select **Offline Restore** to run a restore when the database is offline.

If a restore is run with the database online, other applications can access the database while the restore is in progress. Online restores are available only for partial database restores (and not for full database restores).

- In **Select Backup Images**, click the appropriate full backup image and (if present) associated incremental and delta incremental backup images and then click > to move these images to the **Selection** list (and thereby select the images for restore).

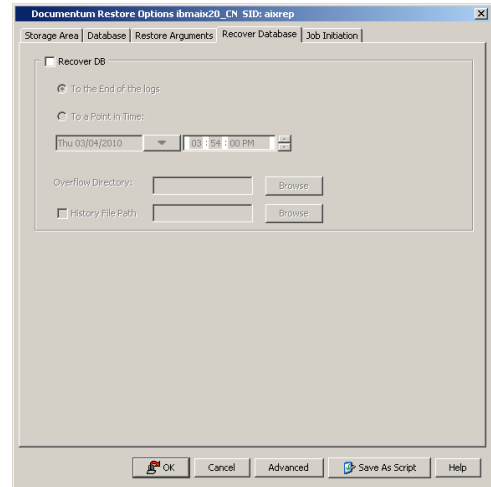


Report files contain the names of the files that have become unlinked during restore operation as a result of a fast reconcile. File names, if specified, must be fully qualified. These files are restored only if the table being restored has DATALINK column type and linked files. To generate this file, you must either enter or browse for and then select the appropriate DB2 Report File path.




8. To restore a DB2 database without selecting backup images, from the **Documentum Restore Options (Recover Database)** tab:

- Select **Recover DB** to recover a DB2 database without selecting backup images from the CommCell Console. Whenever this option is selected, the **Restore Arguments** and **Advanced** tabs are disabled. This option is disabled whenever the options for a partial database restore and/or and out-of-place restore are selected.
- Select **To a Point in Time** to recover the DB2 database based on the state and contents of the database to a specific point in time. To change this time, click in the list as appropriate.



9. If you need to run the restore operation immediately, select **Immediate** from the **Job Initiation** tab.


Go to step 14.



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


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

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



12. From the **Advanced Schedule Options** dialog box, select a specific time zone from which the job schedule time will be based.

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



13. Click **OK** to close the **Schedule Details** window.

14. Click **OK** to close the **Restore Options** window.

You can track the progress of the restore job from the **Job Controller** window.

Once completed, the details of the job can be viewed in the **Job History** window.

VIEWING RESTORE JOB HISTORY

Once a restore job has completed, it may be useful to view specific historical information about the job, such as the following:

- View Restore Items; items in the job that were **Successful**, **Failed**, **Skipped** or **All**.
- View Job Details
- View Events of the restore job.
- View Log files of the restore job.

This information can be viewed in the Job History window, which is accessible at the client computer, /DataAgent, instance, or subclient level in the CommCell Console.

The history provided depends on the entity from which the job history is viewed. For example, viewing job history from the instance level will yield information for the jobs performed for that instance. Conversely, viewing job history from the agent level will yield information for jobs run from all subclients within that instance.

VIEWING RESTORED ITEMS ASSOCIATED WITH A RESTORE JOB

1. From the CommCell Browser, right-click the entity (e.g., instance, client computer, etc.) whose job restore history you want to view, click **View**, and then click **Restore History**.
2. From the **Job History** filter window, select the filter options, if any, that you want to apply, and then click **OK**.
3. From the **Job History** window, right-click the job whose restored items you want to view; click **View Restore Items**, and select from the type of items to

view: **Successful**, **Failed**, **Skipped** or **All**.

4. The **Restored Files** window will display the selected type of restored items for the job.
5. Click **OK**.

[Back to Top](#)

Restore - Documentum iDataAgent (DB2)

Basic | Advanced

TABLE OF CONTENTS

Overview

Component-Level Browse and Restore Operations

- Restoring Storage Areas
- Restoring Full-Text Indexes
- Restoring the Entire Database
- Restoring the Database while Online/Offline
- Restoring History Files
- Restoring Logs
- Reapplying Log Files with Roll-Forward
- Rename one or more Tablespaces with Redirect
- Recovering a Database

Documentum File Browse Operations

- Restoring a Single Document
- Restoring Multi-Versioned Documents
- Restoring a Document to a Different Cabinet

Browse Data

- Selecting Objects From the Browse Window for Restore
- Controlling the Browse Time Interval
- Browsing Data Associated with Specific Jobs
- Global Browse / Restore Options
- Browse Error Reporting
- Browsing and Restoring Using a Specific MediaAgent, Library or Drive Pool
- Browsing and Restoring From Copies

List Media

- Listing Media For a Subclient
- Listing Media For a Backup Set or Instance

Restore by Jobs

Manage Restore Jobs

- Restart Jobs
- Resubmit Jobs
- Control Jobs

Additional Options

OVERVIEW

There will be situations when you might want to restore either the entire database that was lost due to a system failure or hardware malfunction or recover specific portions of the data that was corrupted. Depending on your requirement, you can perform a Component-level Browse and Restore to get back the lost data. In addition, you can also use several alternate restore methods to restore the required data.

Documentum iDataAgent, enables you to restore and recover the following:

- Storage Areas
- Entire database
- Individual data files or table spaces
- Database files, such as log files, history files.
- Full-text Indexes

It is recommended to shutdown the Documentum Repository services before initiating the database restore.

You can also recover the database independently when needed.

Generally, you will want to restore/recover data immediately. However, if you need to restore some large amount of data that is not required immediately, you can schedule it to be restored at some later point of time when either the client computer is not in use or when the network utilization is low. Scheduling the restore operation relieves you of having to manually initiate the operation.

RESTORE DESTINATIONS

By default, the iDataAgent restores data to the client computer from which it originated; this is referred to as an in-place restore. You can also restore the data

to another client computer in the CommCell. Keep in mind the following considerations when performing such restores:

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

The following section enumerates the types of restore destinations that are supported by this iDataAgent.

IN-PLACE RESTORE

When restoring data, the iDataAgent provides the facility to restore the data to the same path/destination on the client computer from which the data originated; this is referred to as an in-place restore.

OUT-OF-PLACE RESTORE

All restore levels provide the facility to restore data to a different directory on the client computer from which the data originated; this is referred to as an out-of-place restore to a different path/destination. You can also restore data to a different client computer in the CommCell using the same path/destination. This can include restoring either individual data files or tablespaces, or the entire database to another computer.

Keep in mind the following considerations when performing out-of-place restores:

- The destination client must reside in the same CommCell as the client whose data was backed up.
- Note that when you perform an out-of-place restore, the restored data assumes the rights (i.e., permissions) of the parent directory.

COMPONENT-LEVEL BROWSE AND RESTORE OPERATIONS

Documentum iDataAgent supports the following types of component-level browse and restore operations.

RESTORING STORAGE AREAS

The Documentum iDataAgent enables you to restore storage areas using the Component-level Browse and Restore method.

RESTORING FULL-TEXT INDEXES

The Documentum iDataAgent enables you to restore full-text indexes using the Component-level Browse and Restore method.

RESTORING THE ENTIRE DATABASE

At times, if the database gets corrupted, you might need to restore the entire database and recover it back to the current state. However, in order to recover a database, you will also need to restore the logs. Both the data and the logs can be restored. You can restore the database in-place or out-of-place.

You can perform a Component-level Browse and Restore to restore the entire database. Minimally, you will need to specify a destination client, the name of the DB2 instance in the **Destination Instance**, the target database name, and target database path.

In the advanced **Restore** tab, you will need to specify options for the archive logs as described in Restoring Logs.

REDIRECT/RENAME DATA DURING RESTORE

One of the advanced restore options available during a restore operation is the ability to redirect datafiles and tablespaces to be restored to a new location. Additionally, datafiles can be renamed during restore. Both of these operations can be performed from the Redirect tab in the Documentum Advanced Restore Options dialog box. Redirecting and/or renaming datafiles and tablespaces as they are restored allows them to be restored without overwriting existing datafiles and tablespaces. This is also useful when restoring and recovering an database to a new host with a different directory structure. While redirecting the tablespaces/datafiles, it is always recommended that you redirect all of the restored tablespaces/datafiles. In other words, restore only the tablespaces/datafiles that need to be redirected.

Select **Switch Database mode for Restore** for putting the database in the appropriate mode before restoring it.

RESTORING THE DATABASE WHILE ONLINE/OFFLINE

The Documentum iDataAgent allows you to restore the database in online or offline mode. When the database is online, other applications can access the database while the restore is in progress. When the database is offline, the restore job has an exclusive connection to the database; any other access to the database is not allowed while the restore is in progress. See step 7. in Component-level Browse and Restore for details.

RESTORING HISTORY FILES

In addition to restoring the entire database, the Documentum iDataAgent also enables you to restore history files.

The database history file is invaluable in a recovery scenario. The history file is individually restorable from any backup image. If the current database is unusable or not available, and the associated recovery history file is damaged or deleted, you can restore the history file. The recovery history file can then be reviewed to provide information on which backup to use to restore the database. The recovery history file provides enough information to recover a database or table space using a backup image. See step 6. in Component-level Browse and Restore for details.

RESTORING LOGS

In addition to restoring data files/ table spaces, the Documentum iDataAgent also enables you to restore the logs. You can also restore the logs based on log time or log serial number.

You can restore the entire logs using the browse and restore method.

To restore the logs:

1. Follow the steps in Component-level Browse and Restore up to step 5.
2. Click **Advanced**.
3. From the **Advanced Restore Options (Log Files)** tab, select **Restore Log Files**.
4. Select either **All Log Files** or **By Point in Time** or **By Log Serial Number** option.
 - o If you select the **By Point in Time** option, specify the point in time in which to restore the Log File.
 - o If you select the **By Log Serial Number** option, specify the restore between two serial numbers.
5. Click **OK** to save the settings and close the **Advanced Restore Options** window.
6. Click **OK** to close the **Restore Options** window and initiate the job.

REAPPLYING LOG FILES WITH ROLL FORWARD

The Documentum iDataAgent allows you to recover the database by reapplying the log files containing transactions that are not part of any database backup image. Whenever log files are restored with the roll-forward option activated, DB2 searches for the following target restore paths per the following order: DB2's own directory path for log file restores; the DB2 Retrieve Path that was specified during the installation of the DB2 iDataAgent; the DB2 Archive Path that was specified during the installation of the DB2 iDataAgent.

To reapply the logs files:

1. Follow the steps in Component-level Browse and Restore up to step 5.
2. Click **Advanced**.
3. From the **Advanced Restore Options (Roll-Forward)** tab, select **Roll-forward (Reapply the Logs)**.
4. Select either **To the End of the logs** or **To a Point in Time** option.
 - o If you select the **To the End of the logs**, all log files will be reapplied.
 - o If you select the **To a Point in Time** option, specify the point in time in which to reapply the Log File.
5. Select an **Alternate Log Path** to provide another location within the DB2 directory for the log files.
6. Select **Leave in Roll-forward Pending State** to specifies whether to leave the database in a roll-forward pending state after the restore. If you select this option, the database will not come up until you roll-forward the database.
7. Click **OK** to save the settings and close the **Advanced Restore Options** window.
8. Click **OK** to close the **Restore Options** window and initiate the job.

DATABASE CONFIGURATION PARAMETERS FOR COMMAND LINE LOG RESTORE

To configure the DB2 database so that log files required during roll-forward when log restore is not selected and the logs are not available on the DB2 database client then the following DB2 database configuration parameters need to be configured using the following commands:

```
db2 update db cfg for <database name> using LOGARCHOPT1''
```

```
CvDctmClientName=<CvDctmClientName>,CvDctmInstanceName=<CvDctmInstanceName>,CvSrcClientName=<CvSrcClientName>,CvClientName=<CvClientName>
```

```
db2 update db cfg for <database name> using VENDOROPT''
```

```
CvDctmClientName=<CvDctmClientName>,CvDctmInstanceName=<CvDctmInstanceName>,CvSrcClientName=<CvSrcClientName>,CvClientName=<CvClientName>
```

where:

<CvDctmClientName> identifies the name of the computer where the Documentum iDataAgent is installed

<CvDctmInstanceName> identifies the name of the Documentum instance configured on client corresponding to <CvDctmClientName>

<CvSrcClientName> identifies the name of the computer where the Documentum iDataAgent is installed

<CvClientName> identifies the name of the computer where the DB2 iDataAgent is installed

<CvInstanceName> identifies the name of the Calypso instance (e.g., Instance001). This should be the name of the DB2 iDataAgent instance associated with the appropriate CommServe.

RENAME ONE OR MORE TABLESPACES WITH REDIRECT

The Documentum iDataAgent allows you to rename one or more table spaces or table space containers upon restore. Renaming these items as they are restored allows them to be restored without overwriting existing table spaces or table space containers.

To rename table spaces with Redirect:

1. Follow the steps in Component-level Browse and Restore up to step 5.
2. Click **Advanced**.
3. From the **Advanced Restore Options (Redirect)** tab, select **Redirect All Table Space to:** and click **Browse** to specify a path.
4. Click **Redirect Table Space Containers** to administer one or more table space containers.
5. Specify names of the **Table Spaces** whose container name can be administered.
6. In **Container Name**, the container name (path) for the table space selected in the **Table Space** field. To take action on this container name, click the name. Once you do this, the **Add**, **Edit**, **Delete**, and **Find and Replace** buttons are activated.
7. Click **OK** to save the settings and close the **Advanced Restore Options** window.
8. Click **OK** to close the **Restore Options** window and initiate the job.

RECOVERING A DATABASE

In certain cases, you might wish to restore the database without selecting backup images.

To recover a database::

1. Follow the steps in Component-level Browse and Restore up to step 8.
2. From the **Recover Database** tab of the **Documentum Restore Options** dialog box and then:
 - o Select **Recover DB** to recover a DB2 database without selecting backup images from the CommCell Console. Whenever this option is selected, the **Restore Arguments** tab is disabled.
 - o Select **To a Point in Time** to recover the DB2 database based on the state and contents of the database to a specific point in time. To change this time, click in the list as appropriate.
 - o Specify an **Overflow Directory** to be search for all of the archived logs during the database recovery.
 - o Specify a History File Path to select the history file that the system will check in order to select the backup image to restore.
3. Click **OK** to save the settings and close the **Advanced Restore Options** window.
4. Click **OK** to close the **Restore Options** window and initiate the job.
5. Click **OK**.

DOCUMENTUM FILE BROWSE OPERATIONS

Before restoring documents, you need to manually make these Documentum jobs inactive using Documentum software:

- dm_DMclean
- dm_DMFileScan

Once the files are restored, re-activate the jobs again.

The Documentum iDataAgent supports the following types of Documentum File Browse restore operations.

RESTORING A SINGLE DOCUMENT

The Documentum iDataAgent allows you to restore a single document. For step-by-step instructions, see Documentum File Browse for more information.

RESTORING MULTI-VERSIONED DOCUMENTS

To browse and restore selected versions of a multi-versioned document to an out-of-place destination and check them into the Documentum Repository:

1. From the CommCell Browser, click a Documentum instance.
2. Right-click a subclient that has been used to back up documents and select **Documentum File Browse**.
3. From the **Browse Options** dialog box, click **OK**.
4. Under the **Name** column, right-click a folder that contains one or more multi-versioned documents and select **Explore**.
5. Expand the folder, right-click a multi-versioned document and select **View All Versions**.
6. Select one or more versions of the document and click **Restore**.
7. The **Restore Options** dialog box appears. Do the following:

- Select the name of the source client computer in the **Destination Client** from the list.
 - Specify a path in **Specify Documentum Destination Folder Path** to perform an out-of-place restore.
 - Click **Unconditional Overwrite**.
 - Click **Create Entities in Repository**.
 - In **Specify Destination Path for Storage Area Files**, enter the file system path on the destination computer to which the files will be restored. Typically, this would be a system path on the Primary Content Server. Click one of the following:
 - **Current Active Version** to check the file in with the same version number that is currently active in Documentum.
 - **Next Minor Version** to check the restored file to the next Minor Version. For example, if the restored document was a Minor version (e.g., 1.2), it will be checked into the next Minor version (e.g., 1.3). If the restored document is a Major version (e.g., 2.0), it will be checked into the next Minor version (e.g., 2.1).
 - **Next Major Version** to check-in the file to the next major version. For example, if the restored document was a Major version (e.g., 2.0), it will check it into the next Major version (e.g., 3.0). If the restored document was a Minor version (e.g., 1.2), it will check it into the next Major version (e.g., 2.0).
8. Click **OK** from the Restore Options dialog box to begin the restore operation.

When restoring a Minor versioned document (e.g., 1.2), it will restore to a Major version (e.g., 2.0). Additionally, if the file exists in Documentum, then the physical file will be restored, but the document check-in will be skipped even when unconditional overwrite is selected.

RESTORING A DOCUMENT TO A DIFFERENT CABINET

To browse and restore selected versions of a document to a different cabinet:

1. From the CommCell Browser, click a Documentum instance.
2. Right-click a subclient that has been used to back up documents and select **Documentum File Browse**.
3. From the **Browse Options** dialog box, click **OK**.
4. Under the **Name** column, right-click a folder that contains one or more multi-versioned documents, and select **Explore**.
5. Expand the folder in the left pane, right-click a document, and select **View All Versions**.
6. Select one or more versions of the document and click **Restore**.
7. The **Restore Options** dialog box appears. Do the following:
 - Click **Unconditional Overwrite**.
 - Unselect **Restore to same folder**.
 - In **Specify Documentum Destination Folder Path**, enter a different `<cabinet_name>/folder` to which the files will be restored.
8. Click **OK** from the Restore Options dialog box to begin the restore operation.

BROWSE DATA

The option to browse the backup data provides the facility to view and if necessary, selectively restore/recover the data that was backed up. This option is especially useful to search for specific data, which includes a specific version of the object.

The browse option can be invoked from the client, agent, instance, backup set, or subclient level depending on the functionality of a given agent. This helps to narrow the search to a specific part of the data.

Depending on the agent, there are several additional options to customize your browse, including:

- Capability to browse the most recent (latest) data.
- Capability to browse data in a specified time range.
- Capability to limit the browse to a specified path.
- Facility to specify the page size of the browse window.
- Ability to browse the image of the data as it existed at the specified browse time.
- Ability to browse from a specified storage policy copy.
- Ability to browse folders/files owned by specific users.

After selecting the necessary browse options, the browse window provides a list of data objects that meets the specified criteria. This window also provides the capability to select multiple or specific data object(s) that you wish to restore/recover.

Note that most of the Unicode / International characters are automatically displayed in the browse window. If these characters are not displayed, make sure that the necessary fonts or other software required to display these characters are installed on the computer in which the CommCell Console is displayed. For example, if you run the CommCell Console as a web-based application, make sure that the necessary fonts or other software required to display the characters are installed on the local computer from which the Console is run.

If necessary, you can also perform the following operations in the browse window, depending on whether the agent supports the operation:

- Search the data using the Find option to restore/recover specific data objects.
- View and restore/recover multiple version(s) of data objects.

Use the following steps to browse data:

1. From the CommCell Browser, right-click the client, agent, backup set or instance, or subclient, click **All Tasks**, and then click the **Browse/Browse Backup Data** option from that level.
2. From the **Browse Options** dialog box, click **OK**.
The Browse window will be displayed.

SELECTING OBJECTS FROM THE BROWSE WINDOW FOR RESTORE

The browse window displays data objects such as files, folders, databases, mailbox, etc. 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 tree.

Note that the window displays only the data that was obtained by a data protection operation. 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.

The buttons in this dialog box include the following:

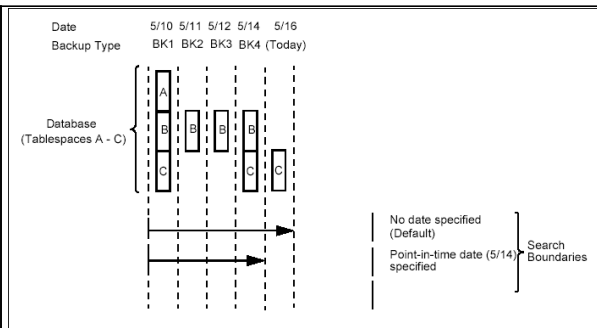
Recover All Selected	Provides access to the Restore Options dialog box which allows you to control how the selected data must be restored.
New Browse	Provides access to the Browse Options dialog box which allows you to re-browse with new browse settings.

CONTROLLING THE BROWSE TIME INTERVAL

The browse operation provides you with the ability to **Browse the Latest Data** or **Browse to Point-in-Time**. The **Browse the Latest Data** option is selected by default. This option returns a list of objects that existed any time and have a backup at the time the command is issued.

You can browse to a point-in-time by specifying the date and time in the **Browse To Point-in-Time** box. The list returned in the Browse window includes all objects which were backed up before the specified date and time.

A table space or data file must have been backed up during the specified browse time span for it to be listed in a browse. If a table space or data file was created and deleted and no backup was performed it will not show up in the browse window.




THE TIME-OF-DAY ELEMENT

The specifications for **Exclude Data Before** options include not only the date, but the time-of-day (i.e., hours and minutes) as well.

Specifying the time is necessary when isolating a backup on a date on which two or more backups occurred. (Note that this condition can occur even if backups are scheduled only once a day. For example, someone may have launched a backup using the Run Immediately option in addition to a scheduled backup. Also, depending on the size of a backup and the time it is scheduled to begin, a backup can start on one date and complete on the next, since the backup need only span 12:00 midnight.)

In determining whether to include a backup in a search, the system uses the time that a backup completes. The **Exclude Data Before** option causes the system to begin its search on the backup that completed after the specified date and time, unless it encounters a full backup first.

The diagram on your right shows the minimum and maximum times that can be given for the **Browse Data Before** and **Exclude Data Before** options in order to define the search boundaries as shown. Notice that the point of delineation is the backup completion time,



2:25am in this case.

BROWSE LATEST DATA

1. From the CommCell Browser, right-click the level whose data you want to browse, click **All Tasks**, and then click the **Browse/Browse Backup Data** option from that level.
2. From the **Browse Options** dialog box, click **OK** to execute the browse using the **Browse the Latest Data** option.

BROWSE DATA BY SPECIFIC TIME RANGE

1. From the CommCell Browser, right-click the level whose data you want to browse, click **All Tasks**, and then click the **Browse/Browse Backup Data** option from that level.
2. From the **Browse Options** dialog box, select the **Specify Browse Time** option and then select the date and time to which you wish to view the data in the browse window.
3. Change the time zone for the specified date and time, if necessary.
4. Click **OK** in the Browse Options dialog box to execute the browse.

BROWSE DATA BETWEEN A SPECIFIED TIME RANGE

1. From the CommCell Browser, right-click the level whose data you want to browse, click **All Tasks**, and then click the **Browse** option from that level.
2. From the **Browse Options** dialog box, select the **Specify Browse Time** option and then select the date and time to which you wish to view the data in the browse window.
3. Change the time zone for the specified date and time, if necessary.
4. Click the **Advanced** button.
5. In the **Advanced Browse Options** dialog box select the **Exclude Data Before** option and then select the date and time from which you wish to view the data in the browse window.
6. Click **OK** in the Advanced Browse Options dialog box.
7. Click **OK** in the Browse Options dialog box to execute the browse.

BROWSING DATA ASSOCIATED WITH SPECIFIC JOBS

You can browse and restore data associated with specific backup jobs. You can do this from the backup history window as follows:

1. From the CommCell Browser, right-click the entity (client computer, iDataAgent, backup set or subclient) whose backup history you want to view, click **View**, and then click **View Backup History**.
2. From the **Backup History** filter window select the filter options, if any, that you want to apply, and then click **OK**.
3. From the **Backup Job History** window, right-click on a job (of a supported agent type) and click **Browse**.

GLOBAL BROWSE / RESTORE OPTIONS

Certain browse options can be set globally, to enable it as the default option. Such options include the following:

- **Displaying Aged Data during Browse**
Aged data which are available on media can be included in the browse operation and subsequent recovery if needed.
Refer to *Accessing Aged Data* documentation for detailed information on this feature.
- **Using Exact Index for Browse**
Use this option to browse data restore and use the index from the specified point-in-time instead of from the last-in-cycle backup. This is useful in situations when you want to restore data from specific media.
- **Disabling Confirmation Message for Restore Jobs**
Confirmation messages for restore jobs can be turned off using this option.
- **Automatically Switching the Copy to Read Data from any Copy**
Use this option to automatically change the copy precedence if the primary copy is unreachable during browse and restore operations.

Follow the steps below to enable these options:

1. From the CommCell Browser, right-click the CommServe, and select **Control Panel**.
2. From the **Control Panel**, select the **Browse/Search/Recovery** option.
3. From the **Browse/Search/Recovery Option** dialog box, select the options you want enabled during the browse and recover operations.

4. Click **OK**.

BROWSE ERROR REPORTING

During a Browse operation, if one of the following error conditions occurs, an accurate problem description will be reported in the **Browse** window. This is extremely useful for troubleshooting.

- MediaAgent Offline
- Index Cache Inaccessible
- Offline Requested Data Not Found

BROWSING AND RESTORING USING A SPECIFIC MEDIAAGENT, LIBRARY OR DRIVE POOL

Data can be restored/recovered from any compatible library and drive type in the CommCell. By default the system automatically identifies and restores/recovers data from any configured library in the CommCell, even if the media is not available in the original library in which the data protection operation was performed. Data Recovery operations use a specific data path - MediaAgent, Library and Drive Pool - to perform the restore operations. When the default options are selected, the system automatically chooses the most appropriate data path. For comprehensive information on choosing an alternate data path, see *Alternate Data Paths*.

In some situations you may want to use another data path to perform a data recovery operation. In such a situation, you can specify the MediaAgent, Library and/or the Drive Pool.

Consider the following examples:

- You may want to use a specific MediaAgent to perform the browse and restore operation instead of the system selected default MediaAgent. For example, the default MediaAgent may be busy and you wish to use another MediaAgent which is idle, or you know the library attached to a specific MediaAgent contains the media associated with the data you wish to restore.
- The default MediaAgent may have a problem accessing the devices (library, drive) and hence you wish to use another MediaAgent sharing the library to perform the browse and restore operation.

This feature is applicable only for tape/optical libraries. The operation will fail if the requested data is not available in the disk library attached to the specified MediaAgent.

CONSIDERATIONS FOR BROWSING DATA

When you perform a browse operation, the system returns the list of files requested by the browse by reading the most recent version of the index in the MediaAgent. If the index is either not available in the MediaAgent's index cache, or not accessible to browse/restore operation then it is restored to the selected MediaAgent from any available MediaAgent. If the MediaAgent used for Browse happens to have visibility to the media then that MediaAgent is used for the browse operation. The same MediaAgent will also get selected in the subsequent restore options. If necessary you can change the MediaAgent in the subsequent restore operation.

BROWSE AND RESTORE USING A SPECIFIC MEDIAAGENT

1. From the CommCell Browser, right-click the agent or instance that contains the data you want to browse, click **All Tasks** and then click **Browse Backup Data**.
2. From the **Browse Options** dialog box, choose the MediaAgent that must be used to perform the browse, from the **Use MediaAgent** drop-down box.
3. If necessary, select the other additional options available in the dialog box.
4. Click **OK** to execute the browse.

RESTORE DATA USING A SPECIFIC MEDIAAGENT, LIBRARY OR DRIVE POOL

1. From the CommCell Browser, right-click the agent or instance that contains the data you want to browse, click **All Tasks** and then click **Browse Backup Data**.
2. From the **Restore Options** dialog box, click **Advanced**.
3. From the **Advanced Restore Options (Data Path)** tab, select the **MediaAgent, Library** or **Drive Pool** that you wish to perform the restore.
4. Click **OK**.

BROWSING AND RESTORING FROM COPIES

By default, when a browse and restore/recover/retrieve operation is requested, the software attempts to browse and restore/recover/retrieve from the storage policy copy with the lowest copy precedence. If the data that you want to browse and restore/recover/retrieve was already pruned from the primary copy, the software searches the other copies of the storage policy in the following order:

1. Lowest copy precedence to highest copy precedence among all synchronous copies.
2. Lowest copy precedence to highest copy precedence among all selective copies (if your agent supports selective copies).

If the data that is to be browsed and restored/recovered/retrieved was secured through multiple storage policies, the software will search for the requested data first from the copy with the lowest copy precedence number to the copy with the highest copy precedence number in the storage policy.

BROWSING AND RESTORING DATA FROM A SPECIFIC COPY

1. From the CommCell Browser, right-click the level whose data you want to browse, click **All Tasks**, and then click the **Browse** option from that level.
2. From the **Browse Options** dialog box, click **Advanced**.
3. In the Advanced Browse Options dialog box select the **Browse from copy precedence** option.
4. Specify the precedence number in **Copy Precedence**.
5. Click **OK** in the Advanced Browse Options dialog box.
6. Click **OK** in the Browse Options dialog box to execute the browse. The Backup Data browse window appears.
7. From the Backup Data browse window, select the objects that you want to restore, and then click **Recover All Selected**. The Restore Options dialog box appears.
8. From the Restore Options dialog box, select the desired restore options. If desired, click **Advanced** to choose additional restore options.
9. Click **OK** from the Restore Options dialog box to begin the restore operation.

For comprehensive information on restoring data from copies, see Browse and Restore/Recover/Retrieve from Copies.

LIST MEDIA

List media option is useful to predict media required for the following operations:

- To restore data associated with a specific backup set, subclient or instance
- To restore the index required to browse data associated with a specific backup set or subclient
- To restore a specific file (or specific files and folders)
- To restore data associated with a specific job

The list media operation can be performed for instances, backup sets, and subclients. The following sections describe each of these methods.

The List Media feature can also be used in a number of other capacities. Refer to the List Media documentation for a complete overview of this feature.

LISTING MEDIA FOR A SUBCLIENT

1. From the CommCell Browser, right-click the subclient for which you wish to list media and then click **List Media**.
2. From the **List Media** dialog box choose one of the following options:
 - Click **Media For the Latest Data** to list media associated with the most recent data protection cycle.
 - Click **Specify Time Range** to list media associated with data protection operations up to the specified date and time range.

Use the **Data Before** box to specify the end date and time.
 - Click **Advanced** and then click **Exclude Data Before** to list media associated with data protection operations after the specified date and time.

Note that you can use the **Specify Browse Time** and **Exclude Data Before** options to list media between a specified date and time range.
3. Click **OK**.

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

LISTING MEDIA FOR A BACKUP SET OR INSTANCE

1. From the CommCell Browser, right-click the backup set or instance for which you wish to list media, click **All Tasks** and then click **Browse**.
2. From the **Browse Options** dialog box, if required, select the following options:
 - Click **Browse the Latest Data** to list media associated with the most recent data protection cycle.
 - 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.
 - 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.

Note that you can use the **Specify Browse Time** and **Exclude Data Before** options to list media between a specified date and time range.
3. Click **List Media**.
4. From the **List Media** dialog box, click **List Media for restore within specified time range** and click **OK**.

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

RESTORE BY JOBS

The Restore By Jobs feature provides the facility to select a specific backup job to be restored. This method of restoring data is considerably faster as it reads continuously on the tape and retrieves the data and does not depend on the indexing subsystem to get the seek offsets on the media.

This feature can be used in different scenarios such as the following:

- To restore point-in-time data associated with a specific backup job, such as full, incremental, differential, etc.
- To restore CommServe DR data for the disaster recovery or for creating a hot-site purposes.
- To restore multiplexed data from the same client.

Consider the following when restoring by jobs:

- Run a restore by jobs whenever you want to restore the entire contents of the job.
- Avoid running restores by jobs for jobs associated with the default subclient. If you do this, the entire contents of the machine will be restored. As such, this may cause problems (e.g., the machine might run out of space) or produce undesirable results (e.g., you might end up restoring operating system files or directories that you really do not want to restore).
- Avoid modifying the contents of any associated subclients since this may result in the retrieval of older data.

Use the following steps to restore data by jobs.

1. From the CommCell Browser, right-click the level whose data you want to browse, click **View History**, and then click the **Backup History**.
2. From the **Backup History Filter** window, select the filter options, if any, that you want to apply and click OK.
3. From the Backup Job History window, right-click the backup job that need to be restored and select **Browse**.
4. From the **Browse Options** dialog box, click **OK** to execute the browse using the **Browse the Latest Data** option.
5. From the **Browse** window, select the instance node in the left pane and then select the required data and logs in the right pane, and click the **Recover All Selected** button at the bottom of the window.
6. Select the desired restore options and click **OK**.

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

RESTART JOBS

Jobs that fail to complete successfully are automatically restarted based on the job restartability configuration set in the Control Panel. This configuration can be changed at any time; however, changes to this configuration will affect all jobs run in the entire CommCell.

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.

Restore jobs for this Agent are re-started from the beginning.

CONFIGURE JOB RESTARTABILITY AT THE COMMSERVE LEVEL

1. In the CommCell Console, click the **Control Panel** icon, then double-click **Job Management**.
2. In the **Job Restarts** tab of the Job Management dialog box, select a job type from the **Job Type** list, and then select the **Restartable** check box.
3. To change the maximum number of times the Job Manager will try to restart a job, select a number in the **Max Restarts** box.
4. To change the time interval between attempts by the Job Manager to restart the job, select a number of minutes in the **Restart Interval (Mins)** box.
5. Click **OK** to save your changes.

CONFIGURE JOB RESTARTABILITY FOR AN INDIVIDUAL JOB

1. From the **Restore Options** dialog box, click **Advanced**, then select the **Job Retry** tab and specify the following as desired:
 - **Total Running Time** - The maximum elapsed time, in hours and minutes, from the time that the job is created.
 - **Number of Retries** - The number of times that Job Manager will attempt to restart the job.
 - **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".
2. Click **OK**.

RESUBMIT JOBS

If a restore job fails to complete successfully, you can resubmit the job without the need to reconfigure the original job's restore options using the **Resubmit Job** feature. When a job is resubmitted, all the original options, restore destinations, and other settings configured for the job remain in tact.

RESUBMIT A RESTORE JOB

1. From the CommCell Browser, right-click a client computer whose data recovery history you want to view, click **View**, then click to view a job history.
2. From the **Job History Filter** dialog box, select **Restore**.
 - o If you want to view more advanced options for restores, from the Job History Filter, select **Restore**, then click **Advanced**.
 - o From the **Data Recovery History Advanced Filter** select the destination client computer of the restores you would like to view, then click **OK**.
3. The system displays the results of the options you selected in the **Data Recovery Job History** window.
4. Right-click on any job, and select **Resubmit**.
5. Select the job options.
6. Click **OK**.

CONTROL 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. Only preemptible jobs can be suspended.
RESUME	Resumes a job and returns the status to Waiting, Pending, Queued, or Running depending on the availability of resources or the state of the operation windows and activity control settings. Restore jobs for this Agent are resumed from the beginning.
KILL	Terminates a job.

These controls can be applied to:

- All jobs in the Job Controller.
- All selected jobs in the Job Controller providing you have the correct security associations at the proper level for each job selected.
- All data protection operations running for a particular client or client/agent.
- All data protection operations running for a particular MediaAgent.

CONTROL A JOB

1. From the Job Controller of the CommCell Console, right-click the job and select **Kill**, **Suspend**, or **Resume** as desired.
 - o When killing a job:

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.
 - o When suspending a job:

The job status may change to **Suspend Pending** for a few moments while the operation completes. The job status then changes to **Suspended**.
 - o When resuming a job:

As the Job Manager attempts to restart the job, the job status changes to **Waiting**, **Pending**, or **Running**.

ADDITIONAL OPTIONS

Several additional options are available to further refine your restore operations. These options are available from the **Advanced Restore Options** window.

To access the **Advanced Restore Options** window:

1. From the CommCell Browser, expand **Client Computers** by double-clicking **Client Computers | Documentum | <instance name>**. The default and other subclients (if available) are displayed on the right-hand windowpane.
 - o To access the Advanced Restore Options for documents, follow the steps in Documentum File Browse up to step 5.
 - o To access the Advanced Restore Options for Documentum Repository components, follow the steps in Component-level Browse and Restore up to step 6.
2. From the **Restore Options** dialog box, click **Advanced**.

The following table describes the additional options:

--	--

OPTION	DESCRIPTION
Data Path Options	<p>The data path options can be accessed from the Data Path tab in the Advanced Restore Options window.</p> <p>Data can be restored to any compatible library and drive type in the CommCell. By default the system automatically identifies the data path for the restore operations as configured in the CommCell. You can change this data path at the restore level by selecting a specific MediaAgent, library or drive pool in the Advanced Restore Options dialog. The data path change can be specified at the restore, schedule and the schedule policy levels. Note that the data path changes specified are only applicable only for the specific restore job/restore schedule.</p> <p>Ensure that the specified MediaAgent and the library or drive pool are available and accessible, failing which the job will go into Pending state; failover protection does not apply here. Also, note that this setting overrides all other settings, including the data path settings at the Storage Policy, Subclient and the alternate data path (Gridstor).</p> <p>In some situations you may want to use another data path to perform a data recovery operation. In such a situation, you can specify the MediaAgent, Library and/or the Drive Pool.</p> <p>Refer to the Change Data Path documentation for a comprehensive overview prior to using this feature.</p>
Startup Options	<p>The start options are accessible from the Startup tab in the Advanced Restore Options window.</p> <p>These options facilitate the startup behavior of the restore job.</p> <p>The Default Priority option dictates the type of job that will be used in determining how the Job Manager will allocate resources for this job.</p> <p>The Change Priority option specifies 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.</p> <p>The 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.</p> <p>Refer to the Job Priority and Priority Precedence documentation for a comprehensive overview prior to using these options.</p>
Run Pre/Post Restore Processes	<p>This option is available from the Pre/Post tab in the Advanced Restore Options window.</p> <p>When you select this option, it allows you to use the restore phases as triggers to start processes. For more information on which agents support Pre/Post Processes for Restore, see Restore Options - Support.</p>
Assign Copy Precedence	<p>This option is available from the Copy Precedence tab in the Advanced Restore Options window.</p> <p>Use this dialog box to assign a particular copy precedence to the copy that you want checked for that particular storage policy.</p> <p>For a more detailed discussion, see Recovering Data From Copies.</p>
Redirect one or more datafiles	<p>These options are available from the Redirect tab in Advanced Restore Options window.</p> <p>You can rename and/or redirect one or more datafiles when the data is restored as well as to redirect table spaces to another location for restore.</p>
Alert	<p>You can configure alerts for the specific restore job from the Alert tab in the Advanced Restore Options window.</p> <p>Refer to the Alerts documentation for comprehensive overview prior to using this feature.</p>

[Back to Top](#)

Restore Data - Documentum iDataAgent - Full System Restore

TABLE OF CONTENTS

Overview

Planning for a Disaster Recovery

Oracle Database Considerations

View Application Version

Perform a Full System Restore

PHASE I. Restore Documentum Components

PHASE II. Restore Database Components

PHASE III. Restore Storage Areas, Full-Text Index Components

OVERVIEW

The difference between a normal restore and a full system restore is the severity of the problem. Normally, if data is lost or removed, it is recovered from the archives using the normal restore procedures. However, when a normal restore operation cannot correct a software and/or hardware corruption problem, some additional changes may be required.

The level of system restore required may be different as described below.

- When the database is corrupted and a restore is required, both the application software and database must be restored. This can be achieved by performing the following steps:
 - Restoring the application by using the file system iDataAgent restore options.
 - Restoring the database files using the restore options in the iDataAgent.
- When the client system (operating system, hardware, hard drives, etc.) is damaged or destroyed, a full system restore may be required.

In the case of disaster recovery, where a full system restore is required, you must rebuild the system to exactly the state as it existed before the problem. In some cases, where the Documentum Repository is corrupted, the Documentum software must be reloaded and the server rebuilt. To resolve a hardware corruption problem, see the appropriate Documentum documentation.

PLANNING FOR A DISASTER RECOVERY

Before an unplanned disaster strikes, do the following:

- Perform regular backups of Documentum and database binaries using the Unix File System iDataAgent.
- Perform regular backups of Documentum configuration files (e.g., `server.ini`, `Docbroker.ini`, `dfc.properties`) and database configuration files using the Unix File System iDataAgent.
- If using a distributed server setup, ensure the correct iDataAgent is installed on the additional computer (e.g., Oracle iDataAgent for the Oracle database or the Unix File System iDataAgent for storage areas). This computer also needs a rebuilding of the Operating System as well as the same patches that were previously installed. Ensure this computer also has access to and can communicate with the CommServe computer. Repeat this for every computer in your distributed environment.
- Using the Documentum iDataAgent, perform a full repository backup that includes the database, storage areas, and full-text indexes.
 - Perform frequent backups of logs and data with control files will also aid in successfully performing a successful disaster recovery procedure.
 - Ensure to exclude database and log files from the backup performed using the Unix File System iDataAgent. This can be achieved by configuring a backup filter in the subclient to exclude these files (e.g., `.dbf`).
- Ensure that the destination client is same as source client while performing a disaster recovery.
- Make sure to restore the configuration files before performing a disaster recovery.
- Once an application has been discovered, the version of the application on which the iDataAgent was installed or upgraded on the client computer is displayed. This can be particularly useful in disaster recovery situations when you need to rebuild the environment.

To view the version of an application,

1. Right-click the instance and select **Properties**.
2. From the **General** tab, view the version of the application in the **Version** field.
3. Click **OK**.

ORACLE DATABASE CONSIDERATIONS

If using an Oracle database, the following considerations are recommended:

- Ensure that the recovery catalog is available on a separate machine. If the recovery catalog is on the same machine, you should have:
 - Exported the user, who is the owner of the recovery catalog using the Oracle **export** command to an external flat file.

- Included this external flat file in your file backup.
- If required, the recovery catalog should be manually resynched using the **Resync Catalog** command.
- Perform the following steps if the recovery catalog is on the same machine that you are rebuilding:
 - Create the recovery catalog database as it existed before the crash.
 - Create the user (who was the owner of the recovery catalog) with the same user privileges that existed for the user.
- Import the user (which was exported to a flat file and was backed up as a part of the file system) using the Oracle **import** command.

PERFORM A FULL SYSTEM RESTORE

A full system restore is the process of restoring all files (including the operating system) and Oracle database files on a client computer after a catastrophic event.

Before You Begin

- Install the same Unix operating system with the same patches that were previously installed.
- Configure the system with the same networking parameters and passwords that were previously set.
- If you are not using DNS, set up peer-to-peer communication with the CommServe. Edit the host file to include the CommServe name.

PHASE I. RESTORE DOCUMENTUM COMPONENTS

▶ On the client computer:

It is recommended to reinstall the Documentum Content Server software and Full-Text Index software (optional), to ensure the integrity of the application software.

1. Install the Documentum iDataAgent software (the Unix File System iDataAgent is automatically installed).
2. If any problems exist with the Documentum iDataAgent, it is recommended that you uninstall and reinstall the software in the same drive as it was previously installed. Use the same installation parameters, passwords, and backup type settings as were previously used.
3. Restore the Documentum parameter files (e.g., `server.ini`, `Docbroker.ini`, `dfc.properties`) which were secured using the Unix File System iDataAgent.
4. Unix File System iDataAgent software. In a clustered environment, you must reinstall to the same virtual node as was being used previously.

PHASE II. RESTORE DATABASE COMPONENTS

▶ On the client computer:

It is recommended to reinstall the Oracle database to ensure the integrity of the application software.

1. Start the database in the **NOMOUNT** mode using the Oracle `Pfile`.
2. Using the Documentum iDataAgent on the CommCell Browser:
 - a. Right-click the Documentum instance that contains the database you want to restore and recover and select **All Tasks**.
 - b. Click **Browse Backup Data**.
3. The **Browse** window is displayed, containing data objects such as database, storage area, full-text index 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 tree.

Note that the window displays only the data that was obtained by a data protection operation. 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.

4. In the Oracle Restore Options (General) dialog, select the following options:
 - a. Select the client computer with the database that you want to restore in the **Destination Client** list (if it is not already selected).
 - b. Click the Control and SP Files tab and ensure that the **Restore Control File** and **SP File** options are selected..
 - c. Click the **Recover** and **Restore Data** tab and ensure that the **Recover** and **Current Time** options are selected.
5. Click on the **Advanced** button and verify the all the required options are selected.
6. Select the appropriate option to begin or schedule the restore, then click **OK**.

Once the job begins, the system will perform a full restore of the selected data, and you can track the status of the job in the Job Controller.

PHASE III. RESTORE STORAGE AREA AND FULL-TEXT INDEX COMPONENTS

▶ On the client computer:

1. Using the Documentum iDataAgent, restore the Storage Areas with inline option. See Restoring Storage Areas for more details.
 2. Using the Documentum iDataAgent, restore the Full-text Indexes and binaries (this is optional).
 3. Backup the associated configuration files along with the Unix File System backup.
 4. Start the Docbroker service.
 5. Start the Documentum Repository service..
 6. Start the FTI Index Server and Index Agents.
 7. Using the Documentum iDataAgent, backup the entire Documentum Repository.
-

Additional Operations - Documentum iDataAgent

TABLE OF CONTENTS

[Audit Trail](#)
[Storage Policy](#)
[Subclient Policy](#)
[Schedule Policy](#)
[Auxiliary Copy](#)
[Operation Window](#)
[Operating System and Application Upgrades](#)
[Uninstalling Components](#)
[Online Help Links](#)

OVERVIEW

In addition to performing backup and restore operations, there are other additional features that can be used to manage and monitor your data and the operations performed on it.

The following sections describes the additional features that can be used with the Documentum iDataAgent.

RELATED REPORTS

Reports are documents you generate from the CommCell console that provide the status of your data management system. You can customize any report you generate by schedule, content and where the report is sent.

CREATING A REPORT

1. From the CommCell Browser, click the **Reports** icon.
2. Click **My Reports**, right click **General** and select **New Report**.
3. Select the report to be generated from the Reports pane in the **Report Selection** dialog box.
4. Select the filter criteria to be include in the report from the corresponding tabs.
5. From the **Output** tab, select the output format and/or save options for the report.
6. Click one of the following: **Run** to generate the report immediately, **Schedule** to schedule the reports for specific times, **Save Template** to save the report filter options into the **My Reports** node, or **Close** to exit the window.

The following reports are applicable for this agent.

- The **Job Schedule Report** filter provides a list of scheduled operations based on the selected filter criteria
- The **Jobs in Storage Policy Copies** Report provides a list of data protection jobs associated with the storage policy copies based on the selected filter criteria.
- The **Backup Job Summary Report** provides a summary of backup jobs for each client.
- The **Calendar Backup Job Summary Report** provides the total amount of backup jobs run (along with their job status) for a specified time period.
- The **Restore Job Summary** provides a summary of restore jobs for each client.
- The **Disaster Recovery Backup Job Summary Report** provides details on all Disaster Recovery backup jobs for a specified time period.

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.

RELATED TOPICS

Additional Operations - Concepts

Provides comprehensive information on additional operations procedures contained on this page.

- Audit Trail
- Storage Policies
- Storage Policy Copies
- Subclient Policies
- Schedule Policy
- Auxiliary Copy
- Operation Window
- Uninstalling Components

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**.

POLICIES

The following policies enable you to plan and schedule the data protection and data recovery operations for your agent.

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. From the CommCell Browser, expand the Policies node, right-click the Storage Policies node, and select **New Storage Policy** from the shortcut menu.
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. Click Next
8. Select the default library for the primary copy from the drop down box and click Next.
9. Select the MediaAgent and from the drop down box and click Next.
10. Enter number of data streams and aging rules and click Next.
11. Click the Browse button and browse to the deduplication store location and click next.
12. The **Review Summary** window is displayed. Confirm your selections and click Finish.

STORAGE POLICY COPIES

Each storage policy has one or more storage policy copies. Storage policy copies enable you to make additional copies of the data. Copies can be created by performing an auxiliary copy operation, or by performing a data protection operation that creates Inline copies.

SUBCLIENT POLICY

Subclient Policies allow you to configure multiple subclients within a CommCell from a centralized template. This alleviates the need to manually configure each file system subclient that shares a similar configuration.

CREATING A SUBCLIENT POLICY

1. From the CommCell Browser, expand the Policies node then right-click the Subclient Policies node and click **New Subclient Policy**.
2. Enter policy name in the **Name** entry field, then select the appropriate file system type from the **iDataAgent** drop-down list. The display pane will show one or more subclient templates listed for your policy.
3. From the Subclient pane assign a Storage Policy for each subclient template by clicking in the Storage Policy column drop down and selecting one from the list.

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. From the CommCell Browser, expand the Policies node, right-click the **Schedule Policies** node, then click **Add**.
2. Type the **Name** of the schedule policy.
3. Select the **Type** of schedule policy from the drop down box.
4. Select the **Agent Type** from the drop down box.
5. Type a description of the schedule policy.
6. Click **Add**.
7. Enter a schedule in the **Schedule Details** dialog box.

8. Click **OK**.
9. Select the **New Schedule Policy - Associations** tab. Select the objects to be associated with the schedule policy.
10. Click **OK**. The schedule policy is displayed in the CommCell Browser.

AUXILIARY COPY

An auxiliary copy operation allows you to create secondary copies of data associated with data protection operations, independent of the original copy. Auxiliary copies have several uses, such as off-site tape storage of scheduled backups.

STARTING AN AUXILIARY COPY

1. From the CommCell Browser, expand the Policies node and right-click the 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.

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.

NEED FOR OPERATION RULE...	APPLY RULE TO...	RULE
Users need uninterrupted access to virtual machines during working hours.	Virtual Server /DataAgent	Disallow data protection operations to run between 8 am and 6 pm, Monday through Friday.

To add an operation rule at the CommServe, Agent, or Subclient level:

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 Rule Details** dialog box:
 - o Enter the name of the rule in the **Name** field.
 - o Select either an administration, data protection (either full or non-full), and/or a data recovery operation from the **Operations** pane.
4. Click **OK** to save the changes.

OPERATING SYSTEM AND APPLICATION UPGRADES

This section contains guidelines and procedures for upgrading the operating system or application(s) on a computer with CommServe, MediaAgent, and/or Agent software.

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), and two methods of upgrading are supported:

- **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.

Whichever path you select to upgrade the OS, all backup data and histories will be preserved and available after the upgrade. 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.

UNINSTALLING COMPONENTS

Once you have uninstalled a component, you will no longer be able to perform data protection operations. The protected data, however, is still available on the media for the length of time expressed in the `Retention Time` defined in the associated Storage Policy Copy. If desired, you can browse the backup data from another active client within the same CommCell or perform data recovery operations from the component to another active client within the same CommCell.

Once you have uninstalled the component software, the corresponding icon in the CommCell Browser will appear dimmed. While this icon exists, most components can still perform data recovery operations. If you delete this icon, all of the component's data is irretrievably lost.

You cannot, however, restore the data to the component's original host computer because the necessary software is no longer present. You will have to reinstall the component software to the original computer in order to perform data recovery operations on such data. See [Reinstalling Uninstalled Agents](#) for comprehensive information on reinstalling the uninstalled agents.

If you want to browse backup data from a client computer whose component software has been removed, simply follow the standard browse procedures.

GENERAL UNINSTALL CONSIDERATIONS

Removing a component is a two-phase process:

- First, you need to uninstall the component software from the client computer hosting the software.
 - If you are unable to uninstall the component from a computer using the recommended procedures, you can use the `Deconfigure Agents` feature to release the license and clean up entries from the CommServe database. However, files and registry entries relating to the component will not be removed from its host computer.
- Next, you must delete the component icon from the CommCell Browser.
 - Once you decide that you want need to browse and/or restore the data after an uninstall or deconfigure, you can manually delete the component from the CommCell Browser. Doing so invalidates the corresponding backup data which may also release backup media to the system. See [Delete the Agent](#) for step-by-step instructions.

Consider the following prior to uninstalling any component:

- Verify that the CommServe is running (except when uninstalling the CommServe software).
- Verify that no jobs are in progress or scheduled to occur while the software is being uninstalled. If jobs are scheduled, either perform the uninstall at another time or disable the jobs on the client computer using the CommCell Console.
- Verify that the CommCell Console and Service Control Manager window are closed.
- If components were not installed in the default location, you will have to manually delete the directory/folders once the uninstall is complete.
- When uninstalling software, you may need to remove additional components and drivers if any had been installed.
- If the **Authentication for Agent Installs** feature has been enabled on the CommCell, users must have Agent Management capabilities to uninstall agents. During the uninstall, users will be prompted for their account information to continue with the uninstall. If the user is not authorized, the uninstall process will abort. For more information, see [Authentication for Agent Installs](#).
- If uninstalling a component from the command line using an uninstall shortcut in the `<software_installation_directory>\Uninstall` folder, the command should not be run with the `<software_installation_directory>\Uninstall` folder open in the command. Instead, the command should be run with the uninstall package included in the command.

For example, the correct command would appear as follows:

```
<software_installation_directory>\Uninstall\<uninstall_package>
```

UNINSTALL SEQUENCE

When uninstalling a component on a Unix or Macintosh computer, the following sequence must be used:

1. Uninstall all application components and updates.
2. Uninstall the MediaAgent (if one exists)
3. Uninstall the Base software

Since the component software depends upon the Base software, you must uninstall the component software first.

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	Documentum Agent Properties (General) Documentum Agent Properties (Version) Documentum Agent Properties (Index) Documentum Agent Properties (Security) Documentum Agent Properties (Activity Control)	
	Instance	Documentum Instance Properties (General) Documentum Instance Properties (Database - Oracle) Documentum Instance Properties (Database - DB2) Documentum Instance Properties (Storage Area and FTI) Documentum Instance Properties (Encryption)	Enter User Name and Password
	Subclient	Documentum Subclient Properties (General) Documentum Subclient Properties (Database - Oracle) Documentum Subclient Properties (Database - DB2) Documentum Subclient Properties (Storage Area) Documentum Subclient Properties (Full Text Index) Documentum Subclient Properties (Activity Control) Documentum Subclient Properties (Storage Device) Documentum Subclient Properties (Security) Documentum Subclient Properties (Encryption)	Add SA Operation Add FTI Operation Data Transfer Option Deduplication Storage Policy Streams
BACKUP	Backup Options	Backup Options	Save As Script Command Line XML Options
	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)	
BROWSE		Browse Options Advanced Browse Options	
RESTORE	Restore Options	Documentum Restore Options (Storage Area and FTI) Documentum Restore Options (Database - Oracle) Documentum Restore Options (Database - DB2) Documentum Restore Options (Restore Arguments - DB2) Documentum Restore Options (Recover Database - DB2) Documentum Restore Options (Job Initiation)	Save As Script Command Line XML Options
	Advanced Restore Options	Advanced Restore Options (General - Oracle) Advanced Restore Options (Copy Precedence) Advanced Restore Options (Data Path) Advanced Restore Options (Encryption) Advanced Restore Options (Pre/Post) Advanced Restore Options (Startup) Advanced Restore Options (Alert) Advanced Restore Options (Redirect) Advanced Restore Options (Options - Oracle) Advanced Restore Options (Restore - Oracle) Advanced Restore Options (Ctrl & SP Files - Oracle) Advanced Restore Options (Recover - Oracle) Advanced Restore Options (Roll-Forward - DB2)	

[Back to Top](#)

Frequently Asked Questions - Documentum

CAN WE PROTECT DATA ON NON-GLOBAL ZONES IF THE FILE SYSTEM /DATAAGENT IS INSTALLED ONLY ON THE GLOBAL ZONE?

Yes. We can protect file system data on non-global zones if the File System iDataAgent is installed only on the global zone. However, in order to enable consistent backups of application specific data on the non-global zones, you will need to install the corresponding application specific iDataAgent on the non-global zone.

HOW DO WE PROTECT DOCUMENTUM DATA ON A NON-GLOBAL ZONE?

In order to enable consistent backups of Documentum data on a non-global zone, you need to install the Documentum iDataAgent on the non-global zone.

ARE ANY DOCUMENTUM JOBS DISABLED DURING BACKUPS?

Documentum cleanup jobs (dm_DMclean) and Documentum filescan jobs (dm_DMFileScan) are disabled during storage area backups. Once the backup job completes, the jobs are enabled again.

IS FULL TEXT INDEX (FTI) SUPPORTED ON WINDOWS AND UNIX STORAGE AREA BACKUPS?

No. Currently, we do not support documentum Full Text Index (FTI) backups when the documentum FTI server is on windows and Unix Storage area backups.

Troubleshoot - Documentum iDataAgent

Backup | Restore

TABLE OF CONTENTS

Recovering Data Associated with Deleted Clients and Storage Policies

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.
 - 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.

[Back to Top](#)